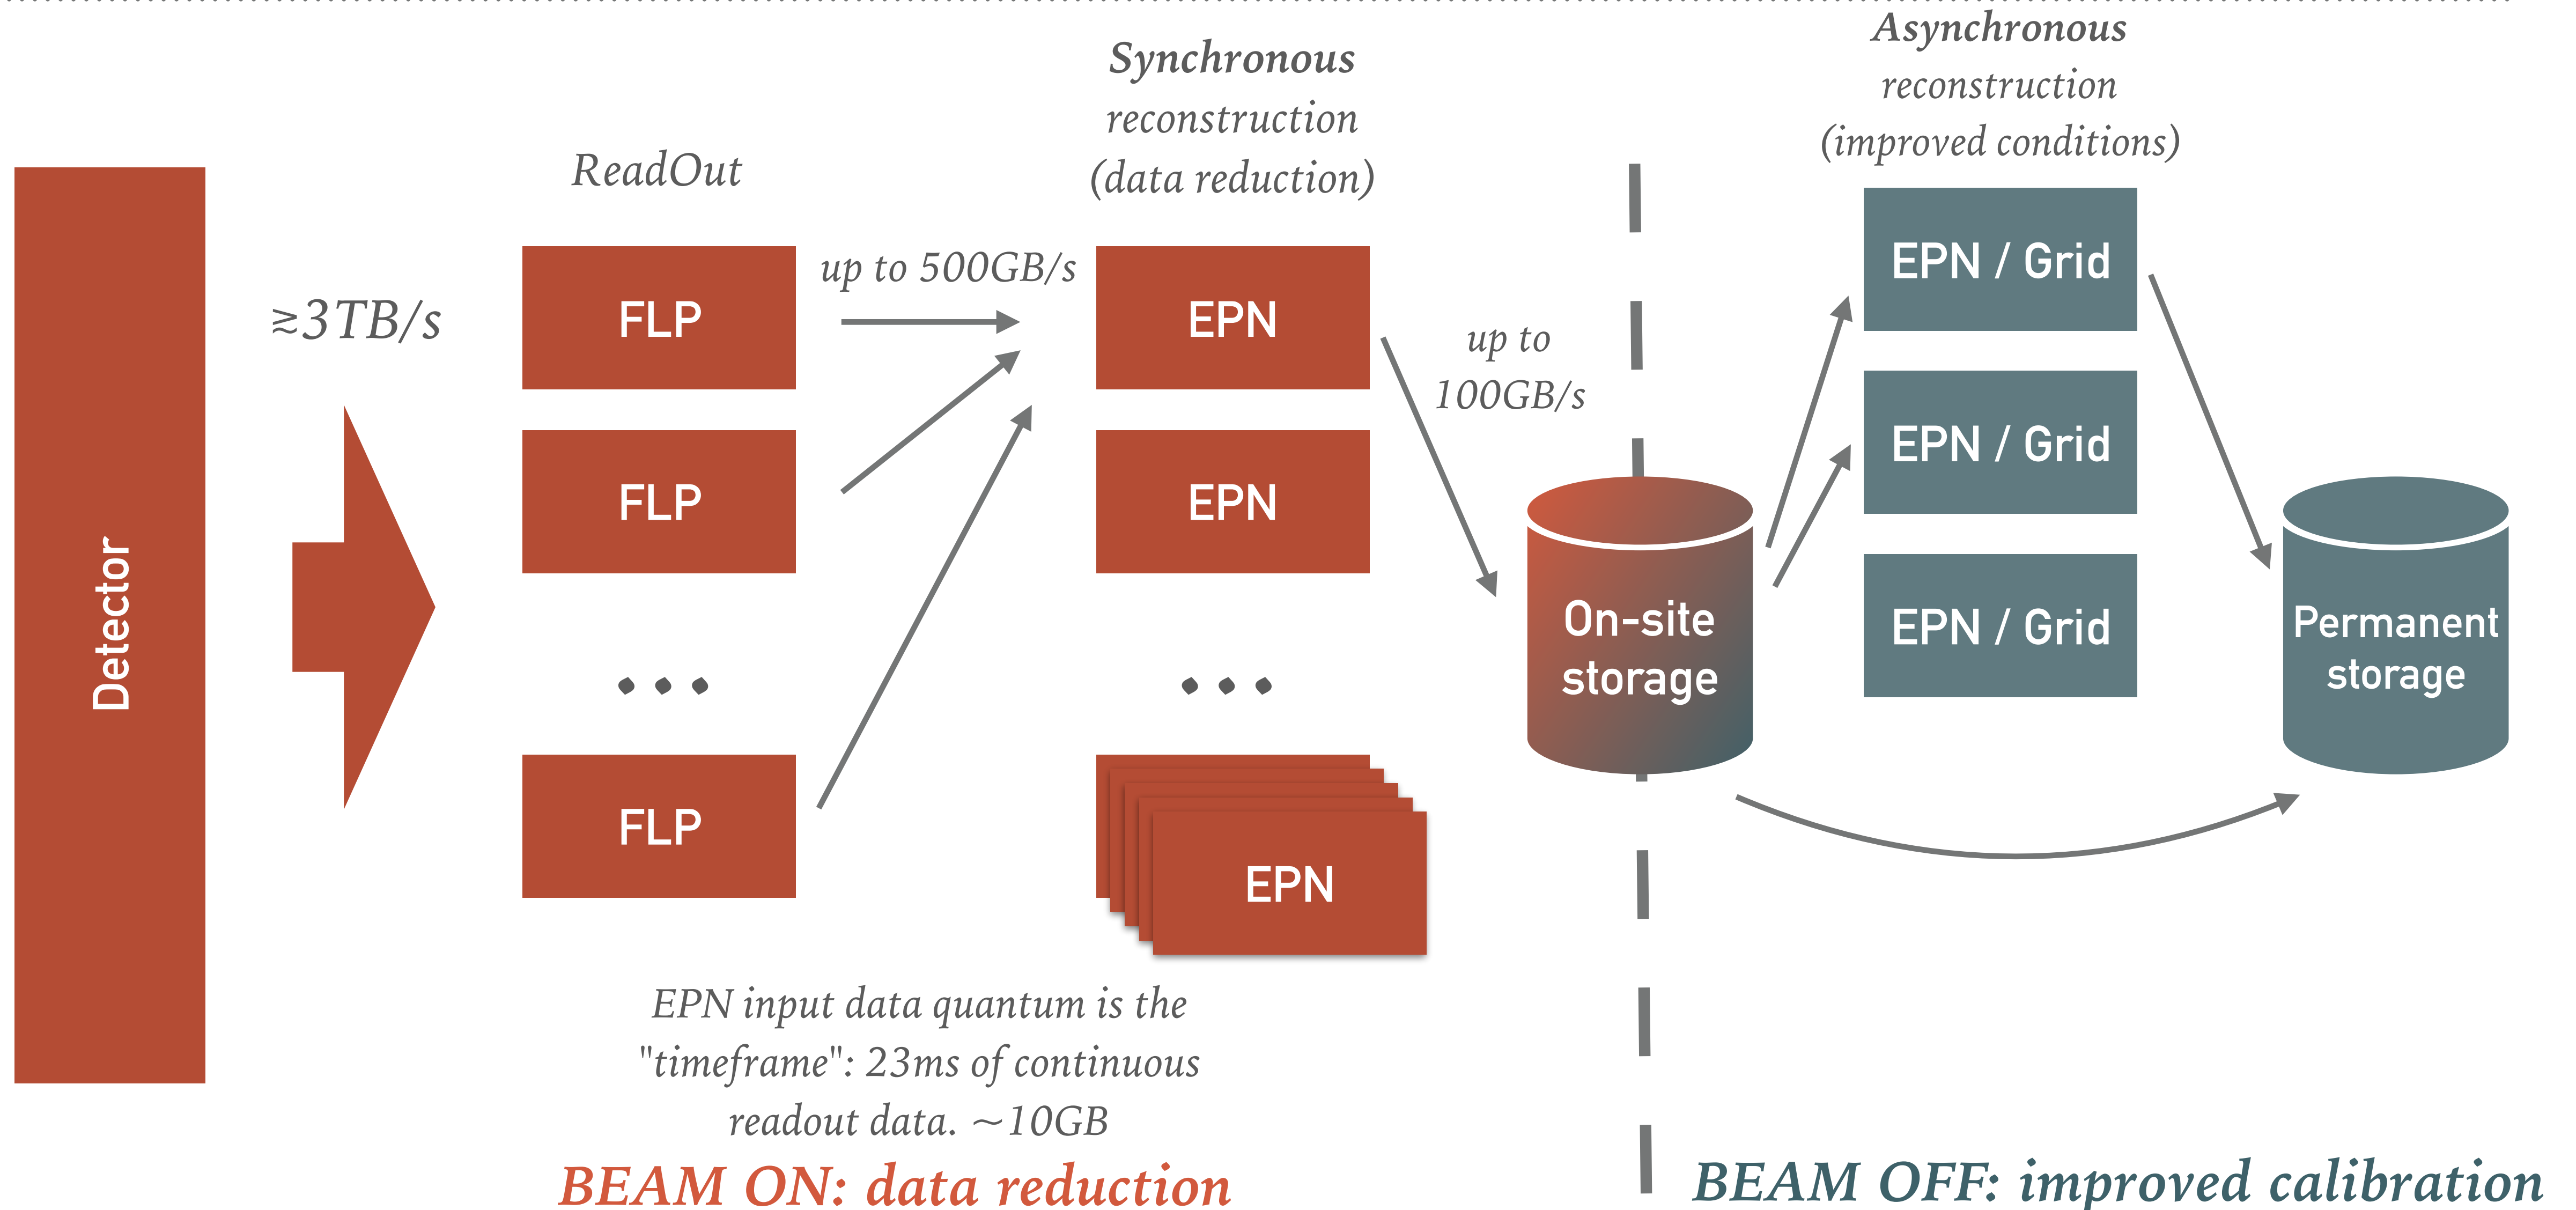


ALICE Software Framework for LHC Run 3

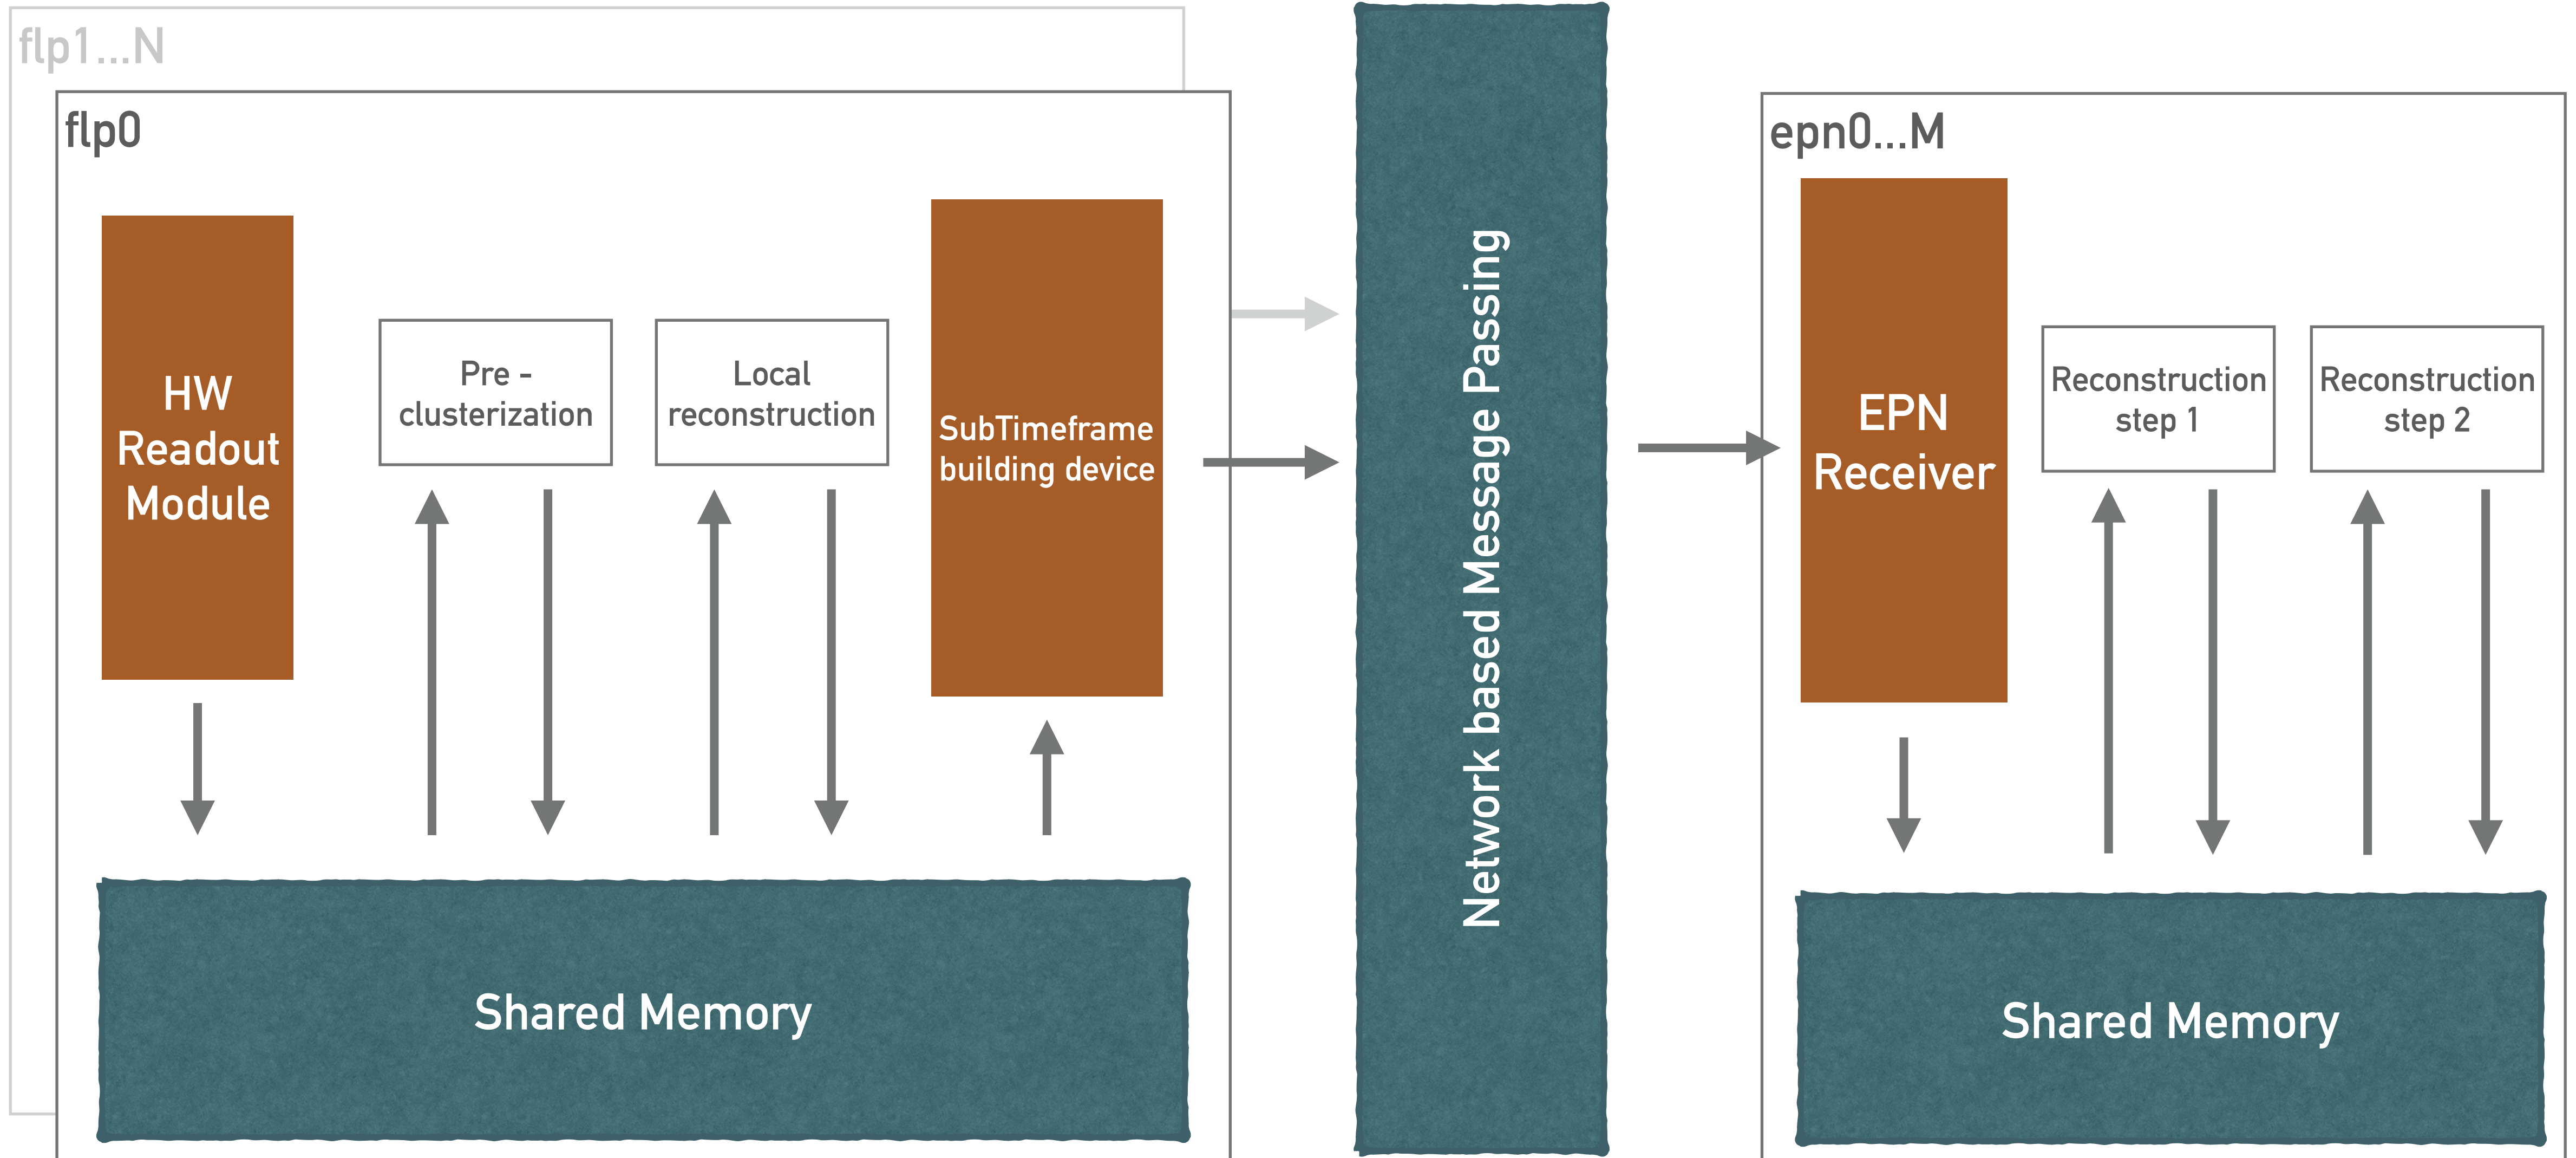


Giulio Eulisse (CERN)

A REMINDER: ALICE IN RUN 3



TRANSPORT LEVEL SYSTEM ARCHITECTURE



ALICE 02 SOFTWARE FRAMEWORK IN ONE SLIDE: FAIRMQ

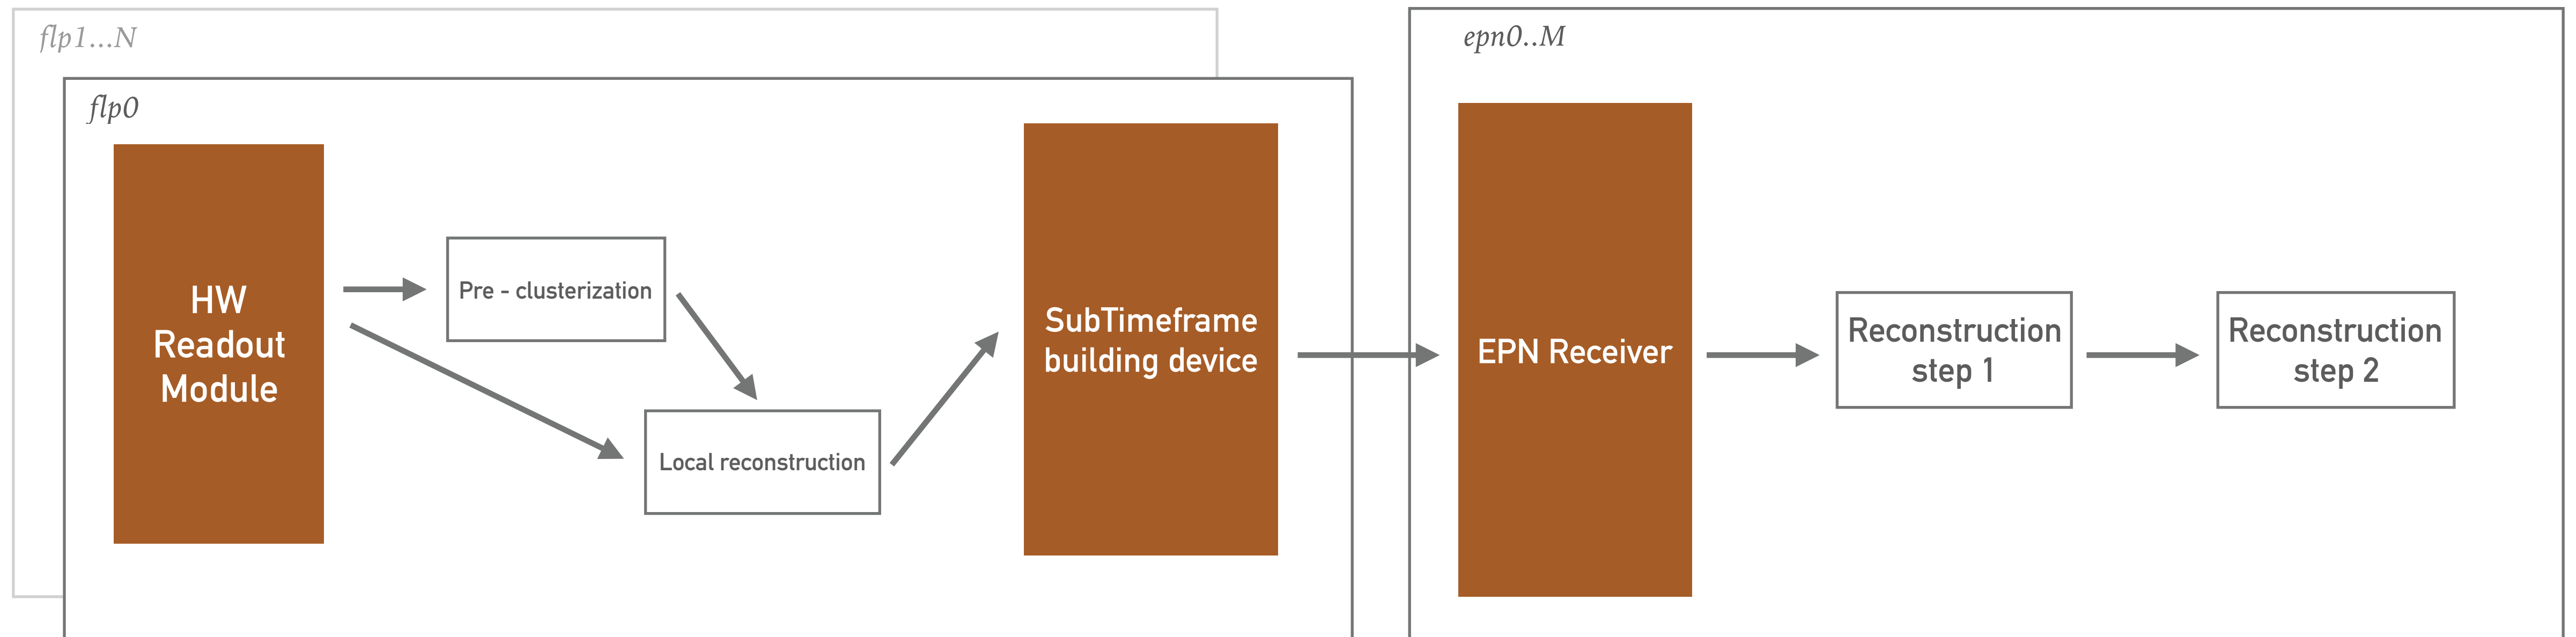
Transport Layer: ALFA / FairMQ¹

- Standalone processes *for deployment flexibility.*
- Message passing *as a parallelism paradigm.*
- Supports both Ethernet and InfiniBand.
- Shared memory *backend for reduced memory usage and improved performance on single node.*

¹*See "ALFA: ALICE-FAIR new message queuing based framework" by Mohammad*

WHY FAIRMQ?

- **Separation of Concerns:** *From the architectural point of view, it allows ALICE to factor out data transport from the system description.*



- **Performant transport:** *collaboration with FAIR experiments and GSI allows sharing of highly skilled engineers to work on the performance critical parts related to transport.*

ALICE 02 SOFTWARE FRAMEWORK IN ONE SLIDE

Data Layer: O2 Data Model


Message passing aware data model. Support for multiple backends:

- ▶ **Simplified, zero-copy format** optimised for performance and direct GPU usage. Useful e.g. for TPC reconstruction on the GPU.
- ▶ **ROOT based serialisation.** Useful for QA and final results.
- ▶ **Apache Arrow based.** Useful as backend of the analysis ntuples and for integration with other tools.

Transport Layer: ALFA / FairMQ¹

- ▶ **Standalone processes** for deployment flexibility.
- ▶ **Message passing as a parallelism paradigm.**
- ▶ **Supports both Ethernet and InfiniBand.**
- ▶ **Shared memory backend** for reduced memory usage and improved performance on single node.

02 DATA MODEL



Messages being exchanged in O2 have a (header, payload) structure where the header describes the contents of the subsequent payload.

- *Origin represents the Detector or Component that first created the message (e.g. TPC)*
- *Description is the data type of the payload (e.g. CLUSTERS),*
- *Subspecification can be used to encode extra information (e.g. TPC sectors)*
- *Timestamp / Timerange indicates the Timeframe it belongs to.*

ALICE O2 SOFTWARE FRAMEWORK IN ONE SLIDE

Data Processing Layer (DPL)

Abstracts away the hiccups of a distributed system, presenting the user a familiar "Data Flow" system.

- *Reactive-like design (push data, don't pull)*
- *Declarative Domain Specific Language for implicit workflow definition.*
- *Integration with the rest of the production system, e.g. Monitoring, Logging, Control.*
- *Laptop mode, including graphical debugging tools.*

Data Layer: O2 Data Model


Message passing aware data model. Support for multiple backends:

- **Simplified, zero-copy format optimised for performance and direct GPU usage.** Useful e.g. for TPC reconstruction on the GPU.
- **ROOT based serialisation.** Useful for QA and final results.
- **Apache Arrow based.** Useful as backend of the analysis ntuples and for integration with other tools.

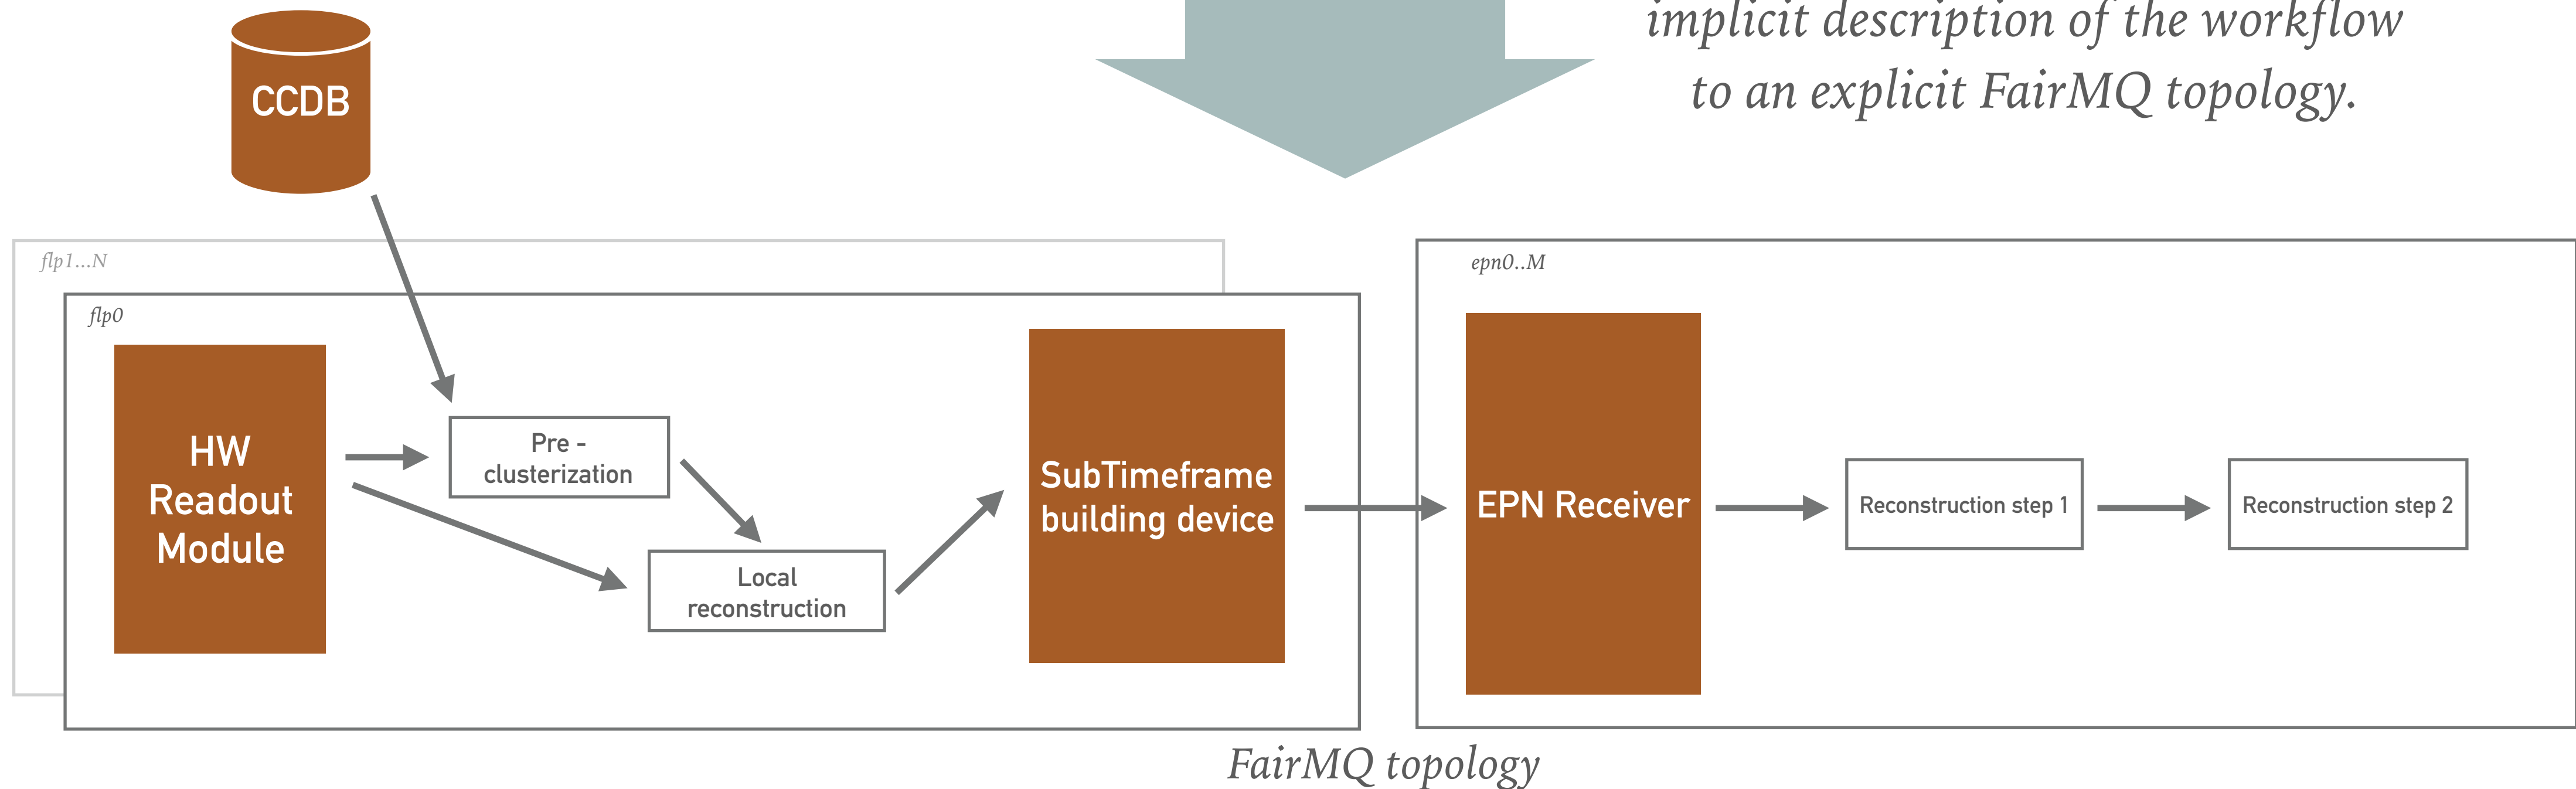
Transport Layer: ALFA / FairMQ¹

- **Standalone processes for deployment flexibility.**
- **Message passing as a parallelism paradigm.**
- **Shared memory backend for reduced memory usage and improved performance.**

DPL IMPLICIT WORKFLOW DEFINITION



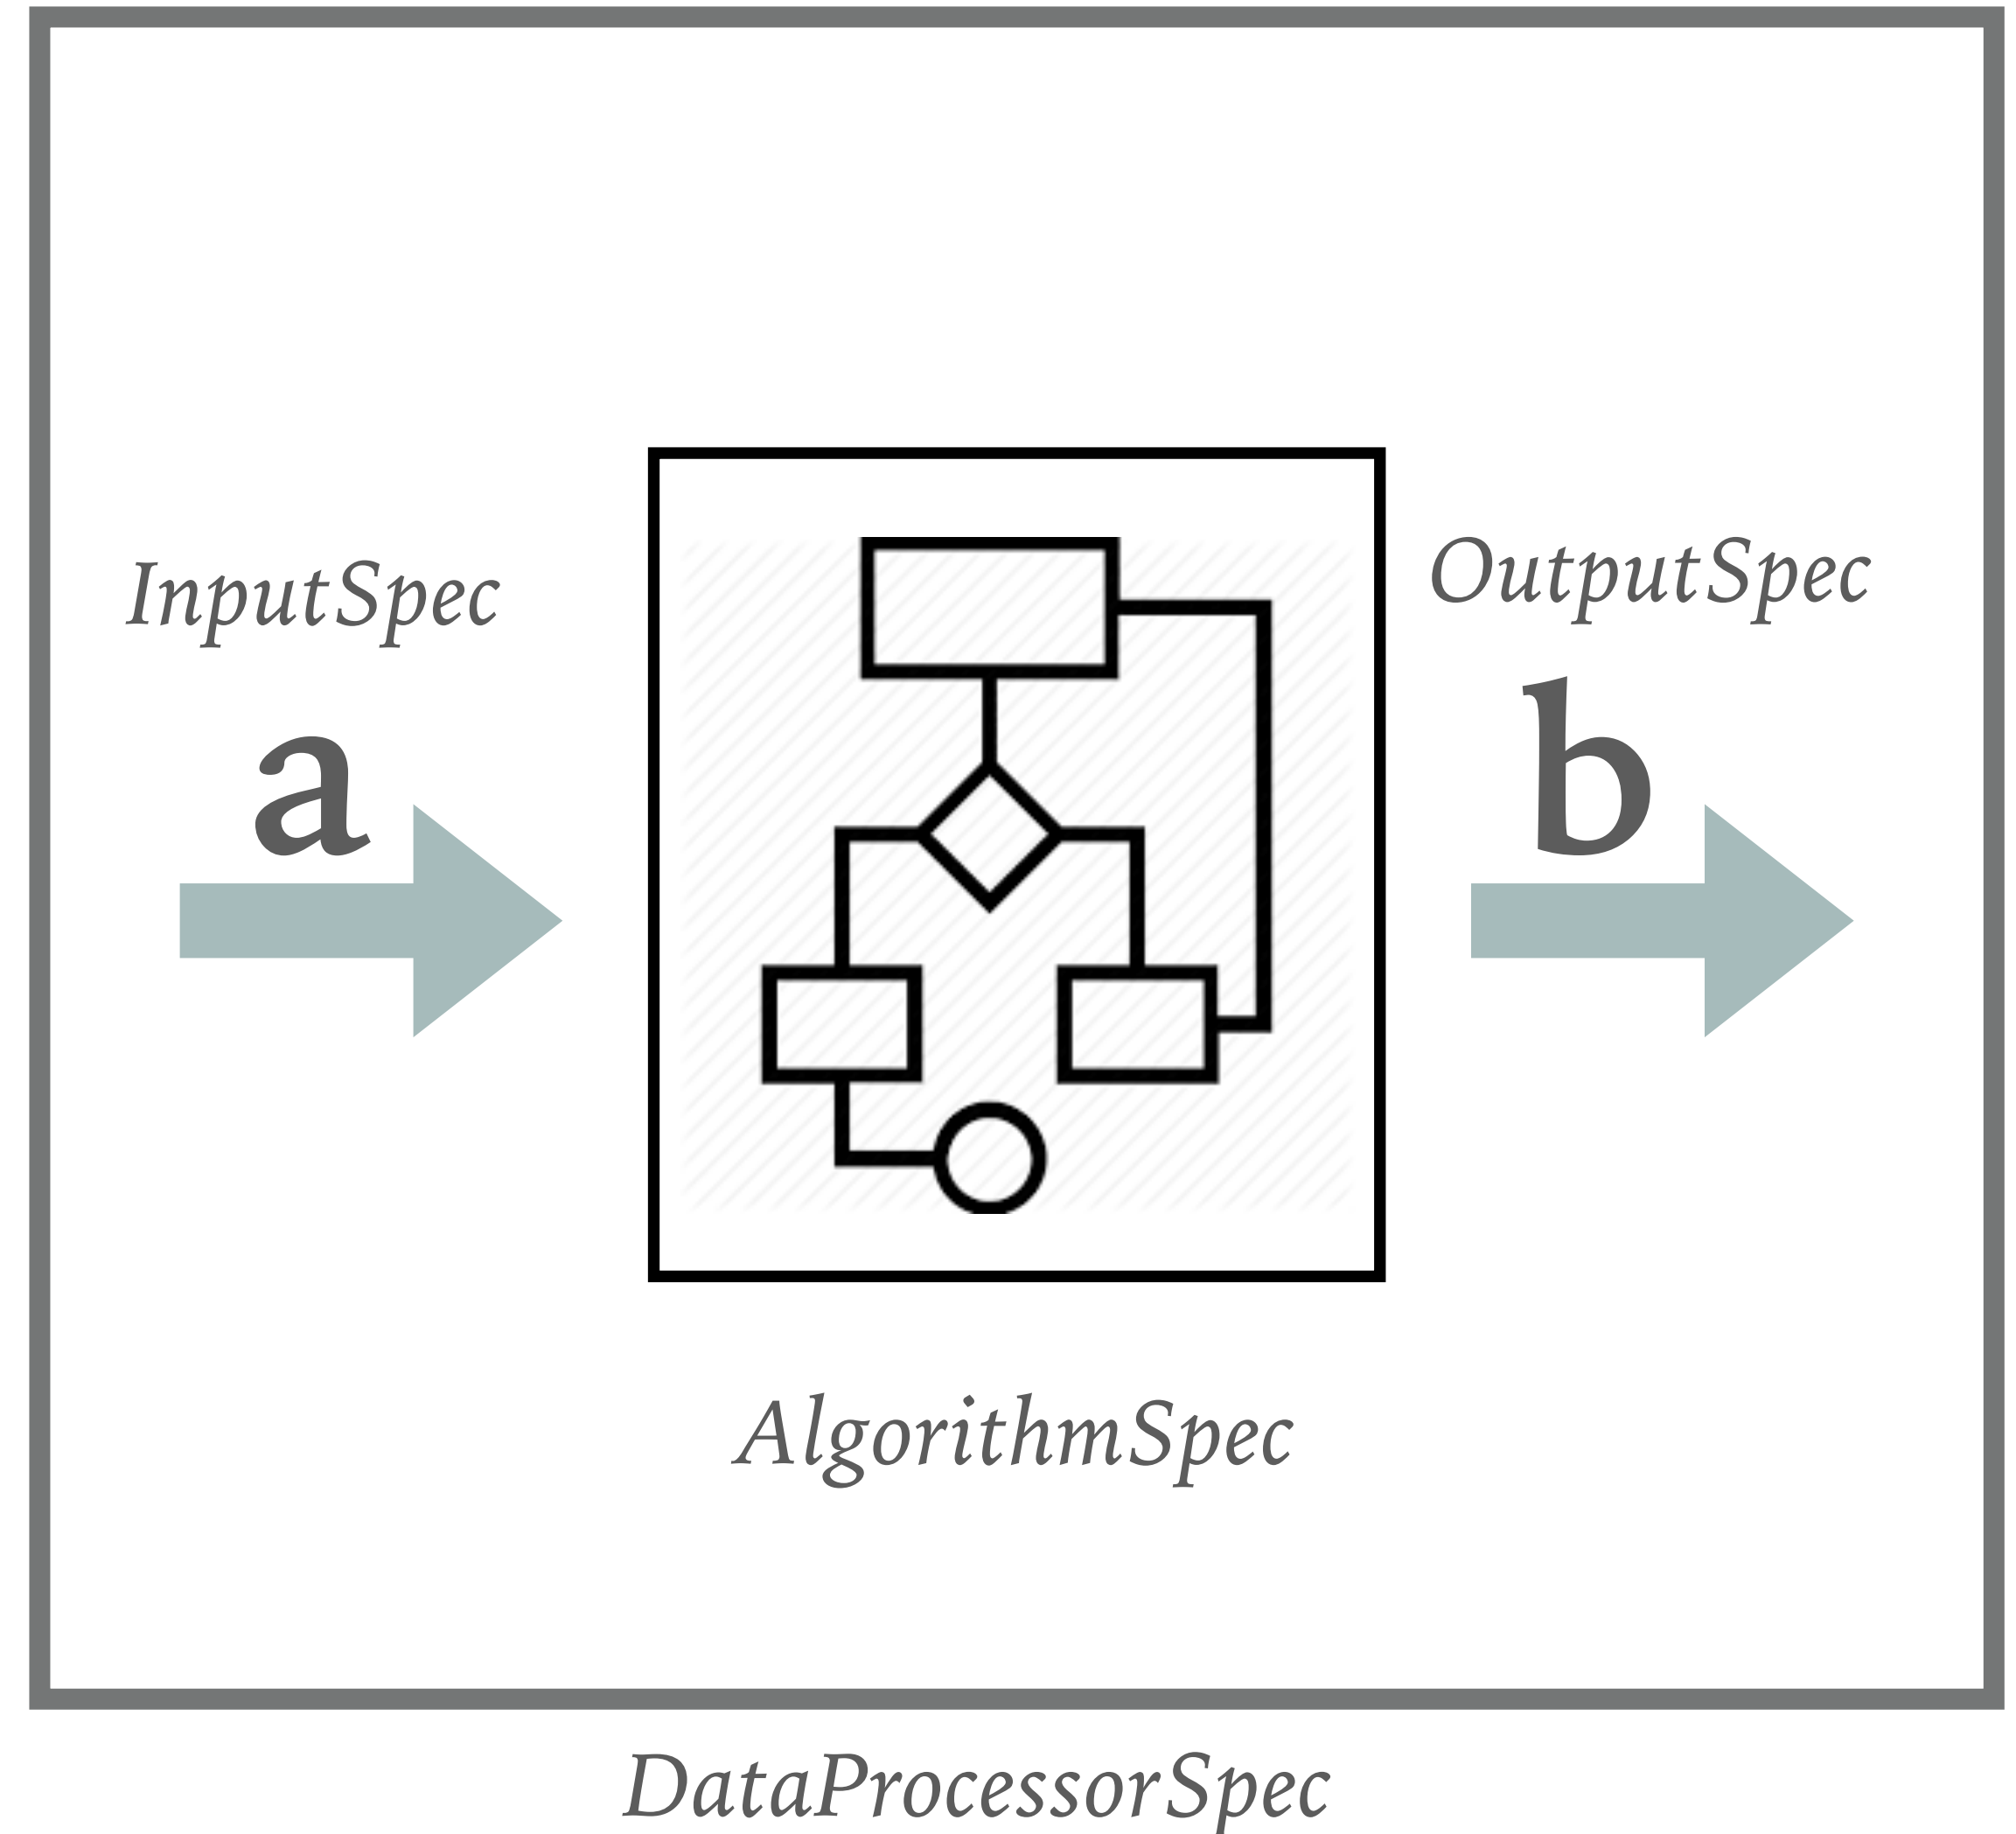
DPL converts a physics oriented implicit description of the workflow to an explicit FairMQ topology.



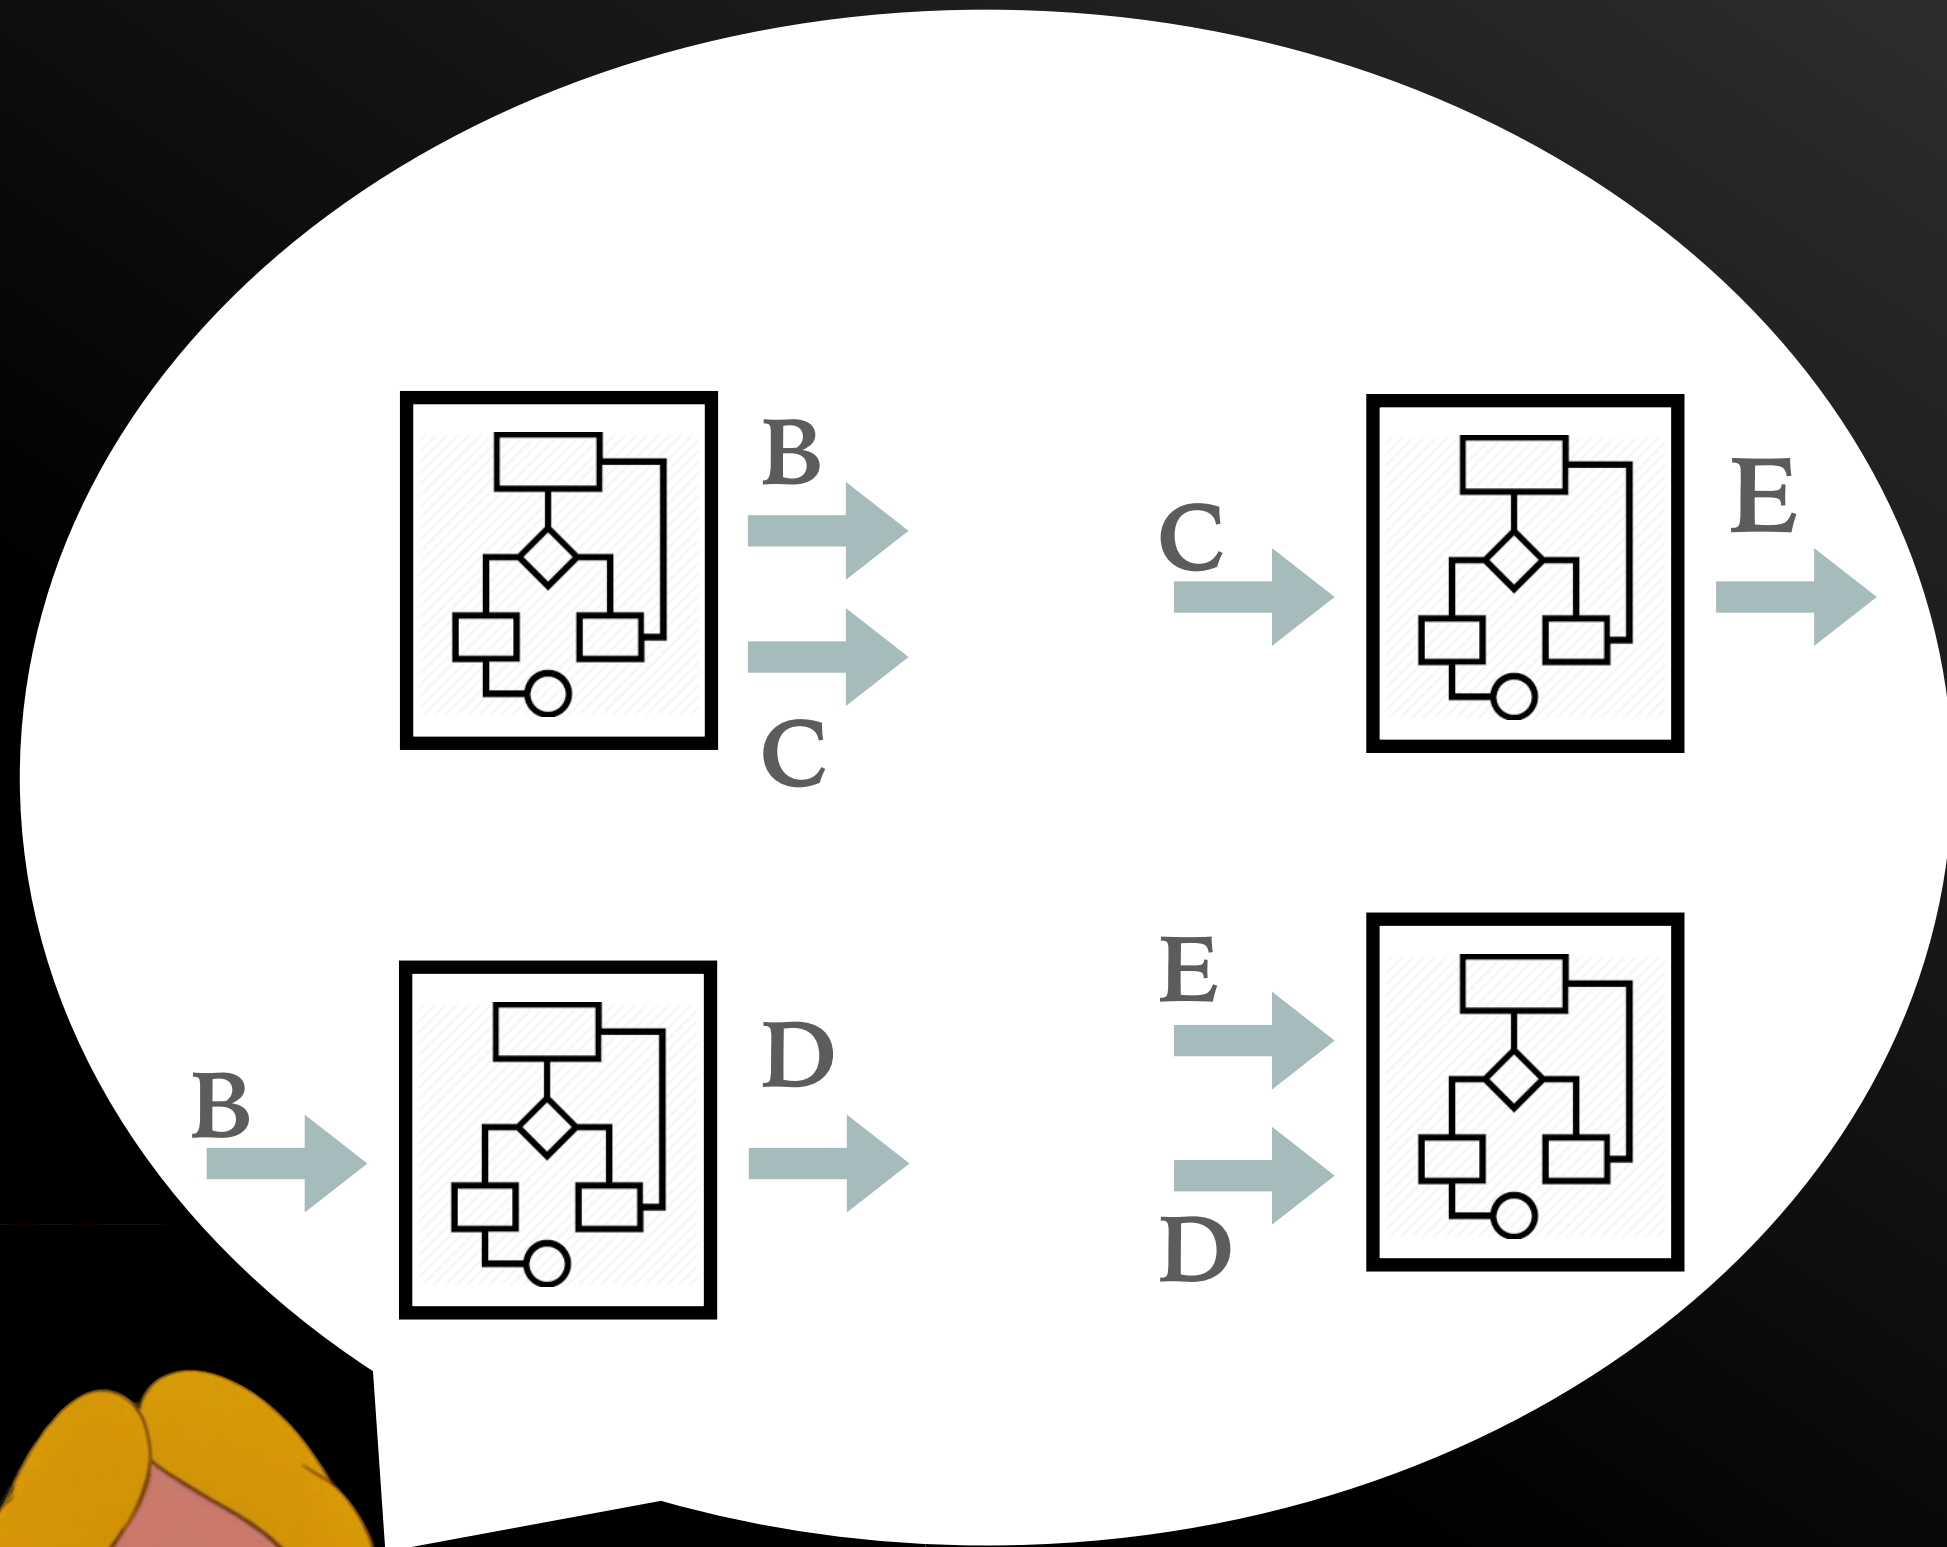
DPL: BUILDING BLOCK

A `DataProcessorSpec` defines a pipeline stage as a building block.

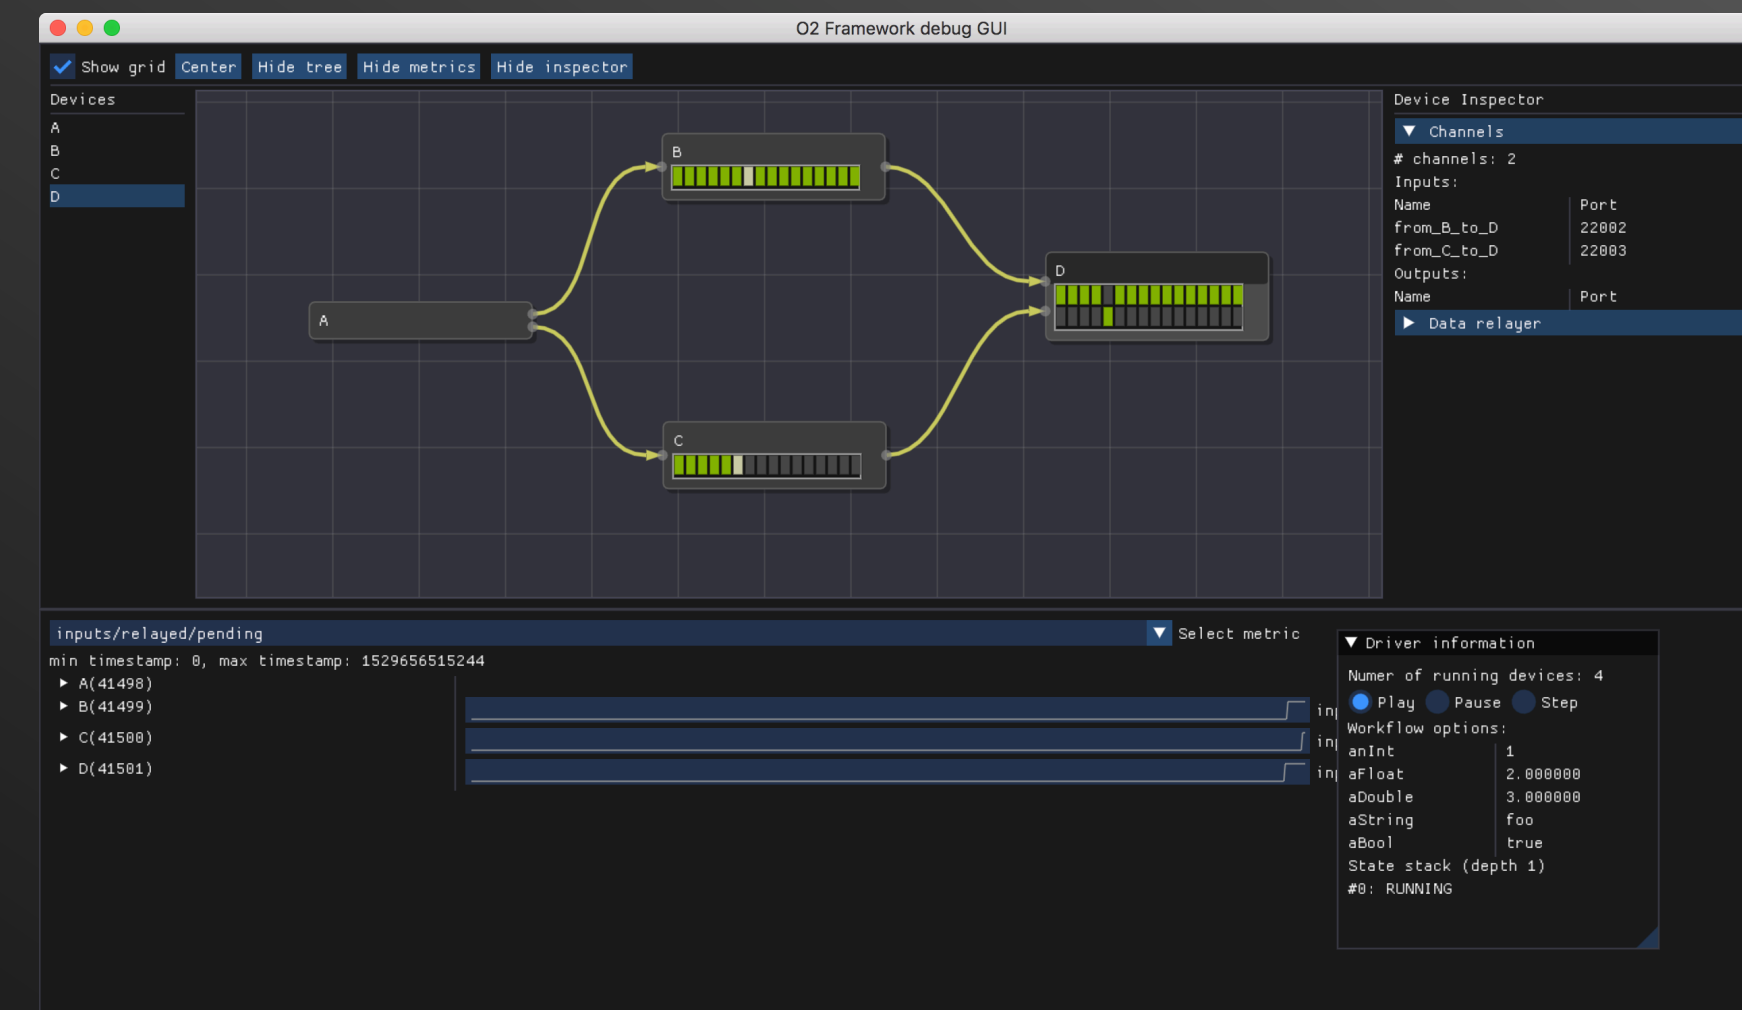
- Specifies inputs and outputs in terms of the O2 Data Model descriptors (*physics types, not sockets and ports*).
- Provide an implementation of how to act on the inputs to produce the output.
- Advanced user can express possible data or time parallelism opportunities.



DATA PROCESSING LAYER: IMPLICIT TOPOLOGY




Data Processing Layer



Show grid
 Center
 Hide tree
 Hide metrics
 Hide inspector

Devices

- A
- B
- C
- D**



Device Inspector

▼ Channels

channels: 2

Inputs:

Name	Port
from_B_to_D	22002
from_C_to_D	22003

Outputs:

Name	Port

▶ Data relayer


Debug GUI

inputs/relayed/pending

▼ Select metric

min timestamp: 0, max timestamp: 1529656515244

- ▶ A(41498)
- ▶ B(41499)
- ▶ C(41500)
- ▶ D(41501)



▼ Driver information

Number of running devices: 4

Play
 Pause
 Step

Workflow options:

aInt	1
aFloat	2.000000
aDouble	3.000000
aString	foo
aBool	true


State stack (depth 1)

#0: RUNNING

Show grid Center Hide tree Hide metrics Hide inspector

Devices

- A
- B
- C
- D



4 FairMQ devices exchanging messages in a diamond topology

Device Inspector

Channels

channels: 2

Inputs:

Name	Port
from_B_to_D	22002
from_C_to_D	22003

Outputs:

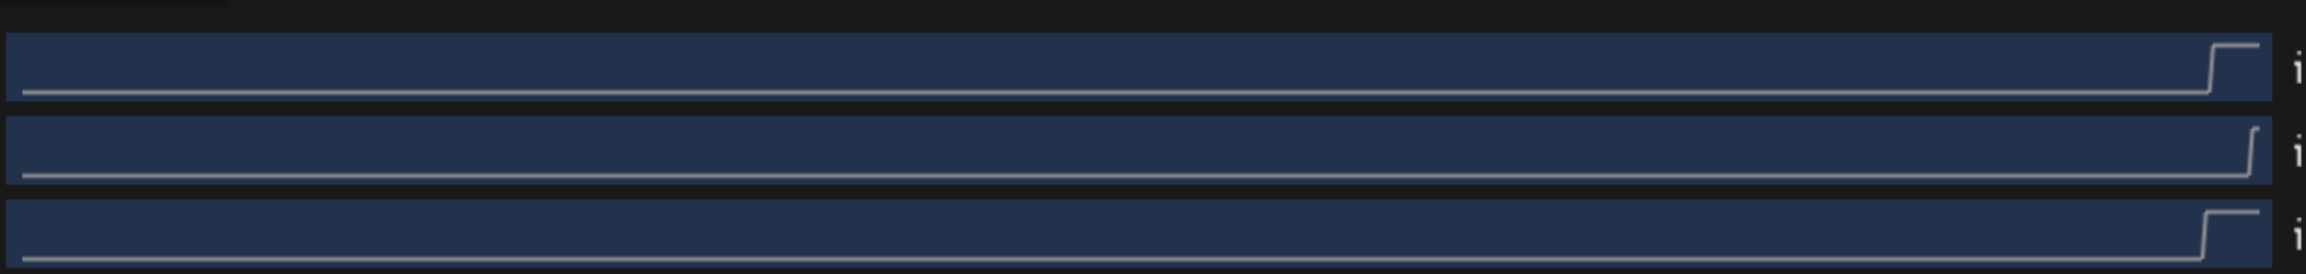
Name	Port
------	------

Data relayer

inputs/relayed

- min timestamp:
- A(41498)
 - B(41499)
 - C(41500)
 - D(41501)

Select metric



Driver information

Number of running devices: 4

Play Pause Step

Workflow options:

aInt	1
aFloat	2.000000
aDouble	3.000000
aString	foo
aBool	true


State stack (depth 1)

#0: RUNNING

Show grid Center Hide tree Hide metrics Hide inspector

Devices

- A
- B
- C
- D



Device Inspector

Channels

channels: 2

Inputs:

Name	Port
from_B_to_D	22002
from_C_to_D	22003

Outputs:

Name	Port
------	------

Data relayer

inputs/relayed/pending

min timestamp: 0, max timestamp: 1529656515244

- ▶ A(41498)
- ▶ B(41499)
- ▶ C(41500)
- ▶ D(41501)

GUI shows state of the various message queues in realtime. Different colors mean different state of data processing.

Select metric

in

in

in

Driver information

Number of running devices: 4

Play Pause Step

Workflow options:

aInt	1
aFloat	2.000000
aDouble	3.000000
aString	foo
aBool	true

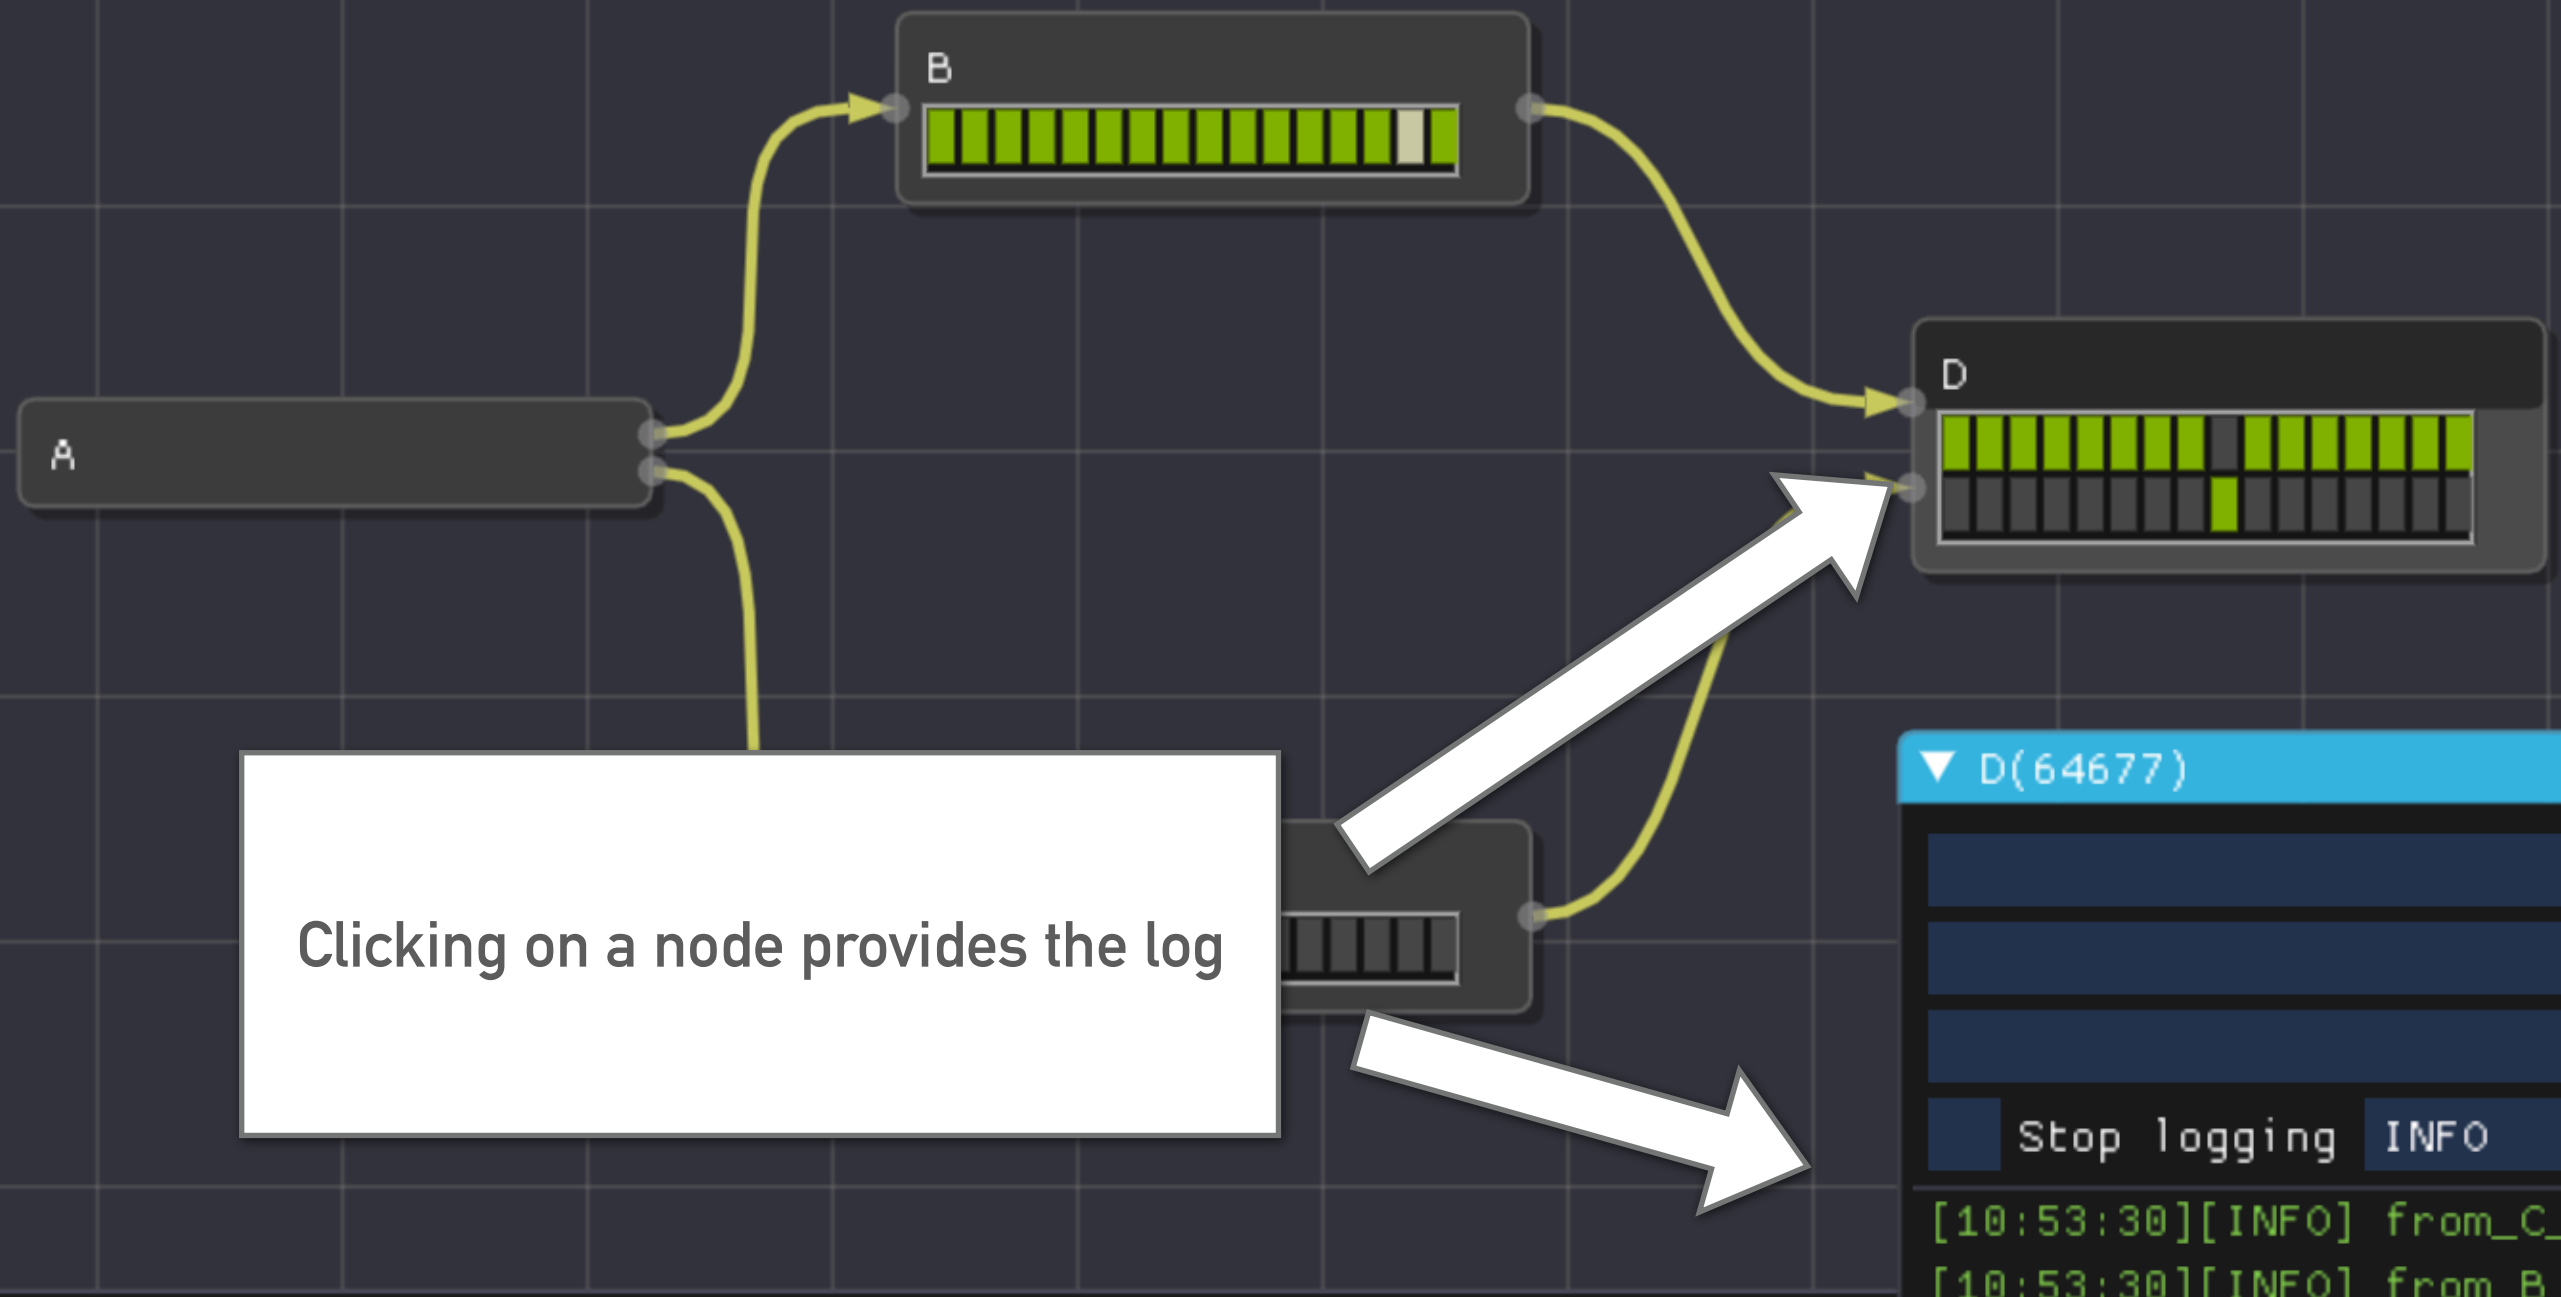
State stack (depth 1)

#0: RUNNING

Show grid Center Hide tree Hide metrics Hide inspector

Devices

- A
- B
- C
- D**



Device Inspector

▼ Channels

channels: 2

Inputs:

Name	Port
from_B_to_D	22002
from_C_to_D	22003

Outputs:

Name	Port
▶ Data relayer	

- ▶ A(64674)
- ▶ B(64675)
- ▶ C(64676)
- ▶ D(64677)

▼ D(64677)

Log filter

Log start trigger

Log stop trigger


Stop logging INFO Log level

```
[10:53:30][INFO] from_C_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:30][INFO] from_B_to_D[0]: in: 0.999001 (0.000131868 MB) out: 0 (0 MB)
[10:53:31][INFO] from_C_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:31][INFO] from_B_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:32][INFO] from_C_to_D[0]: in: 1 (0.000132 MB) out: 0 (0 MB)
[10:53:32][INFO] from_B_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:33][INFO] from_C_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:33][INFO] from_B_to_D[0]: in: 1 (0.000132 MB) out: 0 (0 MB)
[10:53:34][INFO] from_C_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:34][INFO] from_B_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:35][INFO] from_C_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:35][INFO] from_B_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:36][INFO] from_C_to_D[0]: in: 0 (0 MB) out: 0 (0 MB)
[10:53:36][INFO] from_B_to_D[0]: in: 1 (0.000132 MB) out: 0 (0 MB)
[10:53:37][INFO] from_C_to_D[0]: in: 0.995025 (0.000131343 MB) out: 0 (0 MB)
[10:53:37][INFO] from B to D[0]: in: 1.99005 (0.000262687 MB) out: 0 (0 MB)
```

Workflow options:

Show grid Center Hide tree Hide metrics Hide inspector

- Devices
- A
- B
- C
- D



Device Inspector

Channels

channels: 2

Inputs:

Name	Port
from_A_to_C	22001

Outputs:

Name	Port
from_C_to_D	22003

Driver information

Numer of running devices: 4

● Play ● Pause ● Step

Workflow options:

anInt	1
aFloat	2.000000
aDouble	3.000000
aString	foo
aBool	true

State stack (depth 1)

#0: RUNNING

An embedded metrics viewer provides in GUI feedback on DPL & user defined metrics. Multiple backends supported, including InfluxDB (i.e. for ALICE data taking) and Monalisa (Grid deployments).



dpl/stateful_process_count lines

min timestamp: 1531126299592, max timestamp: 1531126385662





```

1 #include "Framework/runDataProcessing.h"
2
3 using namespace o2::framework;
4
5 AlgorithmSpec simplePipe(std::string const &what) {
6     return AlgorithmSpec{ [what](ProcessingContext& ctx) {
7         auto bData = ctx.outputs().make<int>(OutputRef{what}, 1);
8     } };
9 }
10
11 WorkflowSpec defineDataProcessing(ConfigContext const&specs) {
12     return WorkflowSpec{
13         {"A", Inputs{}, {OutputSpec{"a1"}, "TST", "A1"}, OutputSpec{"a2"}, "TST", "A2"}},
14         AlgorithmSpec{
15             [](ProcessingContext &ctx) {
16                 auto aData = ctx.outputs().make<int>(OutputRef{ "a1" }, 1);
17                 auto bData = ctx.outputs().make<int>(OutputRef{ "a2" }, 1);
18             }
19         },
20     },
21     {"B", {InputSpec{"x", "TST", "A1"}}, {OutputSpec{"b1"}, "TST", "B1"}, simplePipe("b1")},
22     {"C", {InputSpec{"x", "TST", "A2"}}, {OutputSpec{"c1"}, "TST", "C1"}, simplePipe("c1")},
23     {"D", {InputSpec{"b", "TST", "B1"}, InputSpec{"c", "TST", "C1"}}, Outputs{},
24         AlgorithmSpec{[](ProcessingContext &ctx) {}}
25     },
26 };
27 }

```

The previous example (GUI included) requires 27 user's SLOC.


DATA PROCESSING LAYER: HOW



DataProcessorSpec

```
DataProcessorSpec{
  "A",
  Inputs{
    InputSpec{"a", "TPC", "CLUSTERS"}
  },
  Outputs{
    OutputSpec{"b", "TPC", "TRACKS"}
  },
  AlgorithmSpec{
    [](ProcessingContext &ctx) {
      auto track = ctx.outputs().make<Track>(OutputRef{ "b" }, 1);
    }
  }
}
```

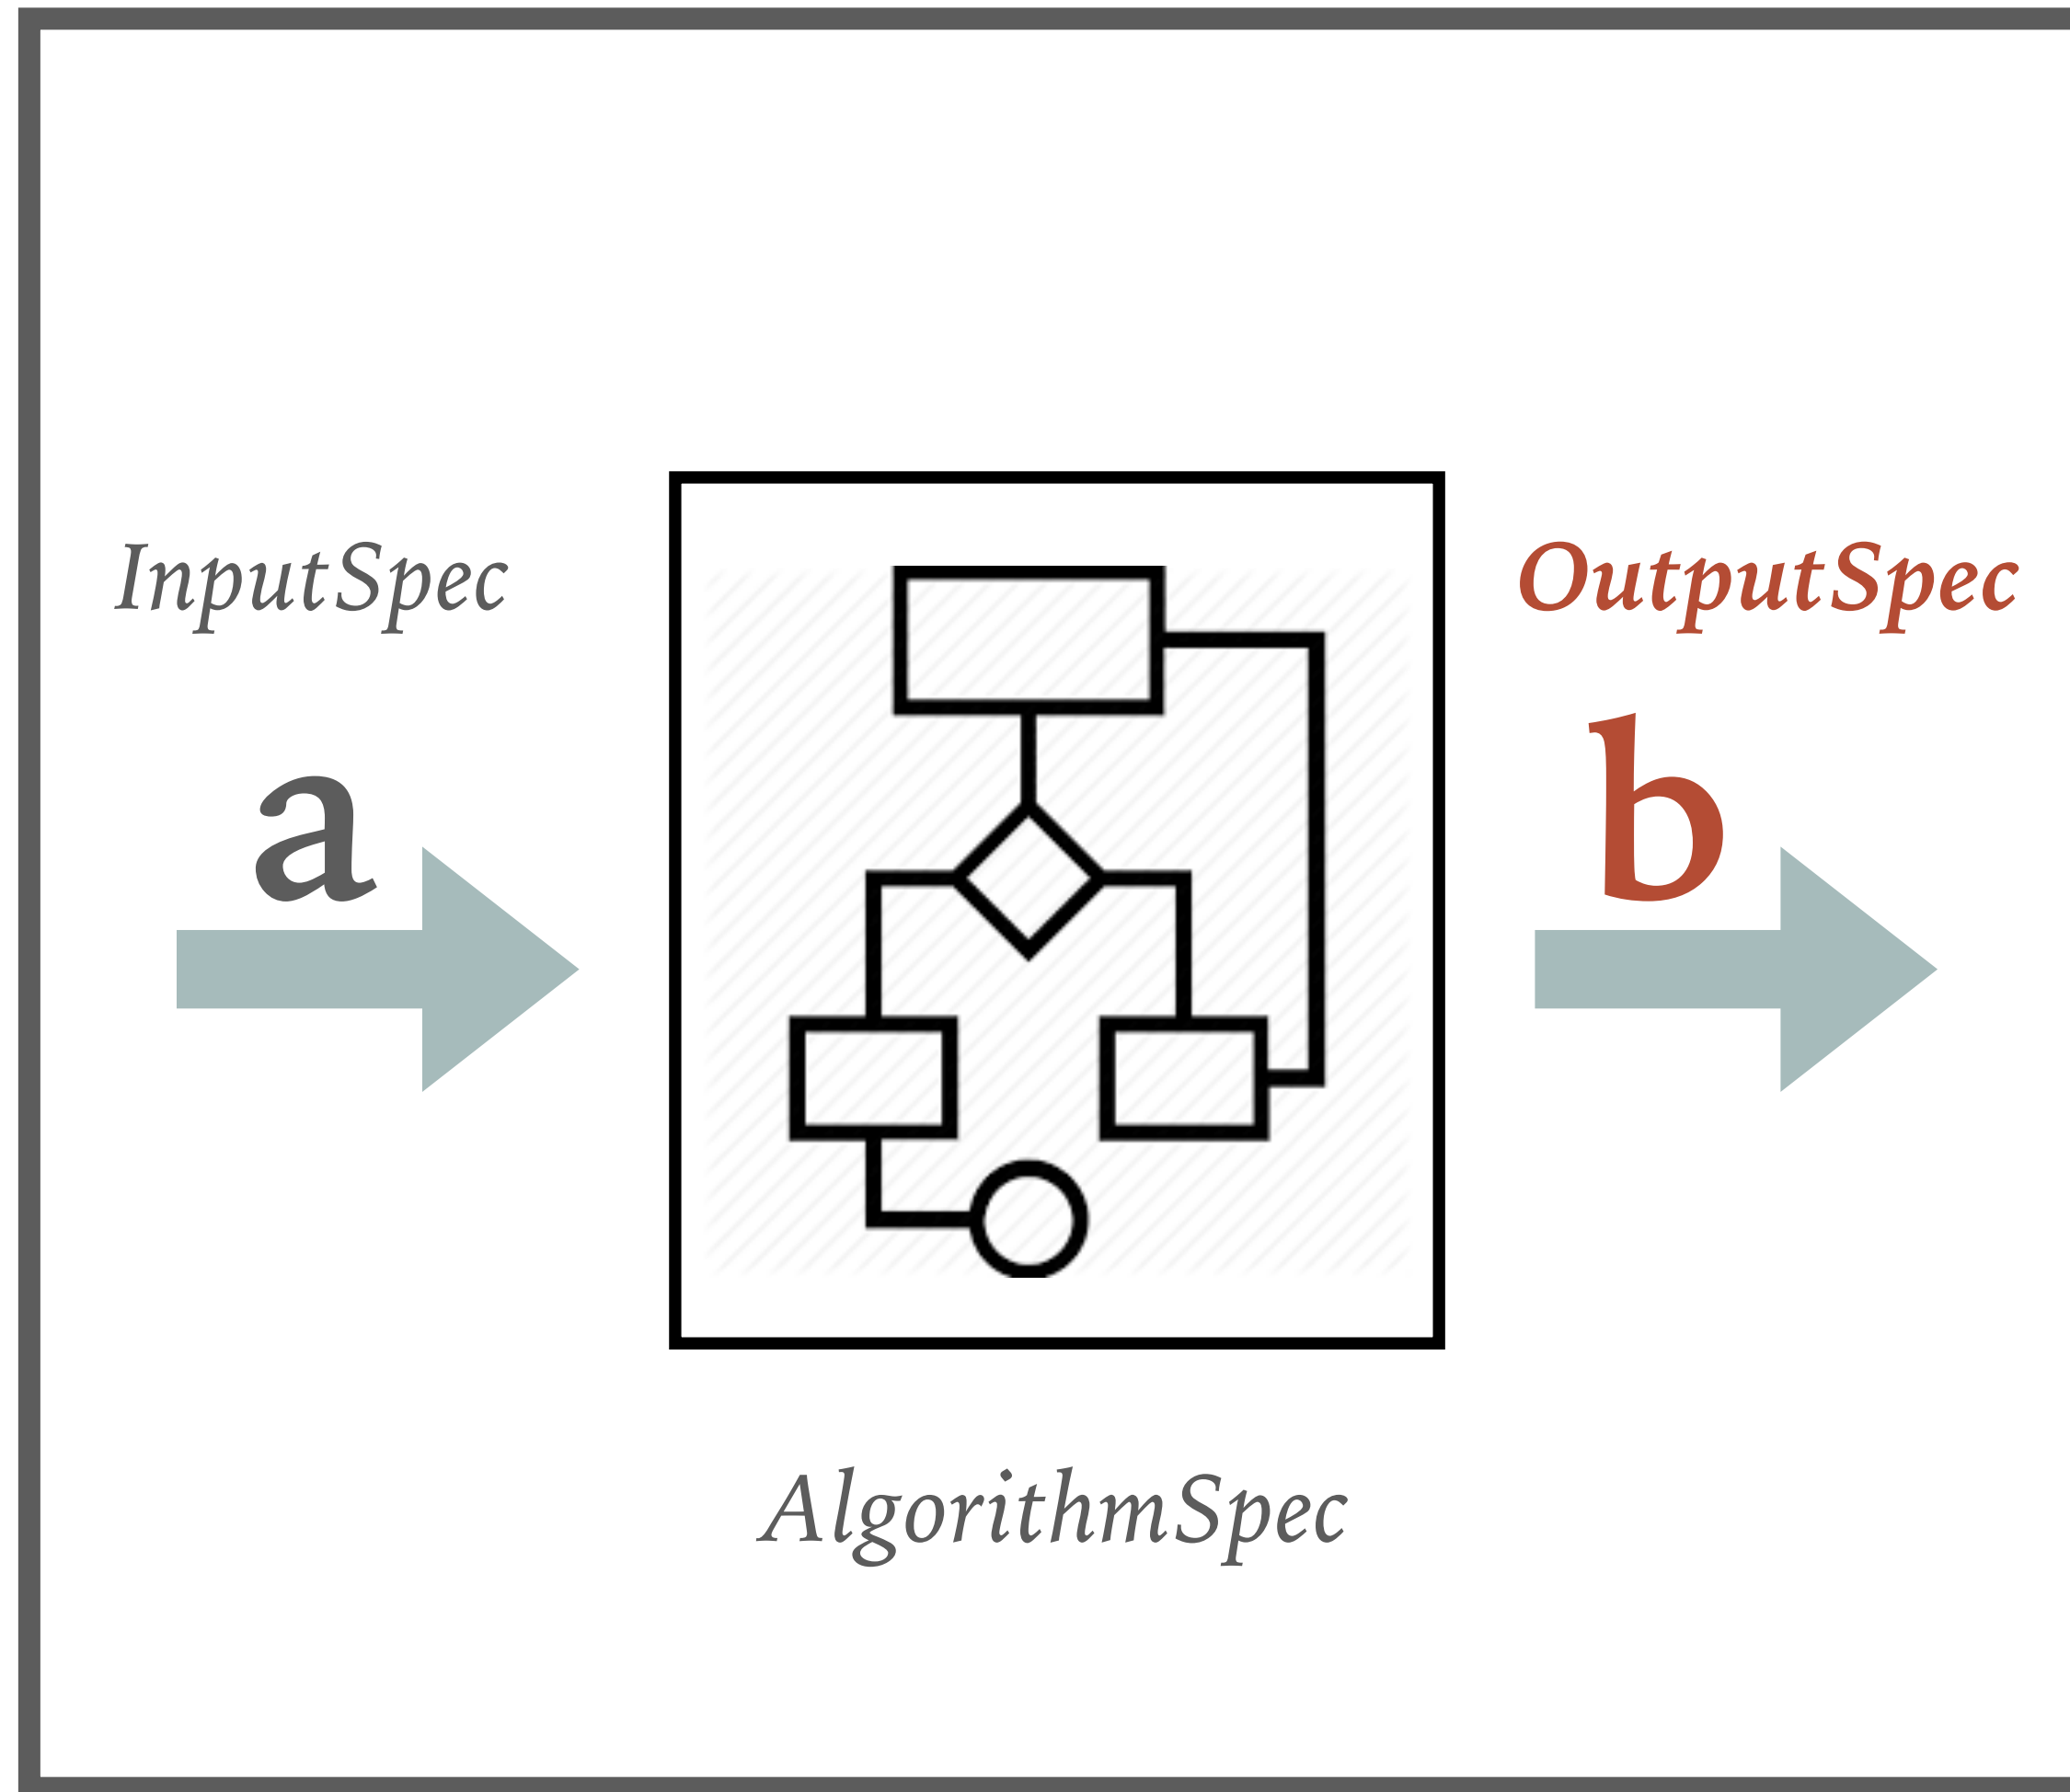
DATA PROCESSING LAYER: HOW



DataProcessorSpec


```
DataProcessorSpec{
  "A",
  Inputs{
    InputSpec{"a", "TPC", "CLUSTERS"}
  },
  Outputs{
    OutputSpec{"b", "TPC", "TRACKS"}
  },
  AlgorithmSpec{
    [](ProcessingContext &ctx) {
      auto track = ctx.outputs().make<Track>(OutputRef{ "b" }, 1);
    }
  }
}
```

DATA PROCESSING LAYER: HOW



```
DataProcessorSpec{
  "A",
  Inputs{
    InputSpec{"a", "TPC", "CLUSTERS"}
  },
  Outputs{
    OutputSpec{"b", "TPC", "TRACKS"}
  },
  AlgorithmSpec{
    [](ProcessingContext &ctx) {
      auto track = ctx.outputs().make<Track>(OutputRef{ "b" }, 1);
    }
  }
}
```


DATA PROCESSING LAYER: HOW




```
DataProcessorSpec{
  "A",
  Inputs{
    InputSpec{"a", "TPC", "CLUSTERS"}
  },
  Outputs{
    OutputSpec{"b", "TPC", "TRACKS"}
  },
  AlgorithmSpec{
    [](ProcessingContext &ctx) {
      auto track = ctx.outputs().make<Track>(OutputRef{ "b" }, 1);
    }
  }
}
```

REACTIVE DESIGN


Data is described as pushed through the pipeline.



REACTIVE DESIGN

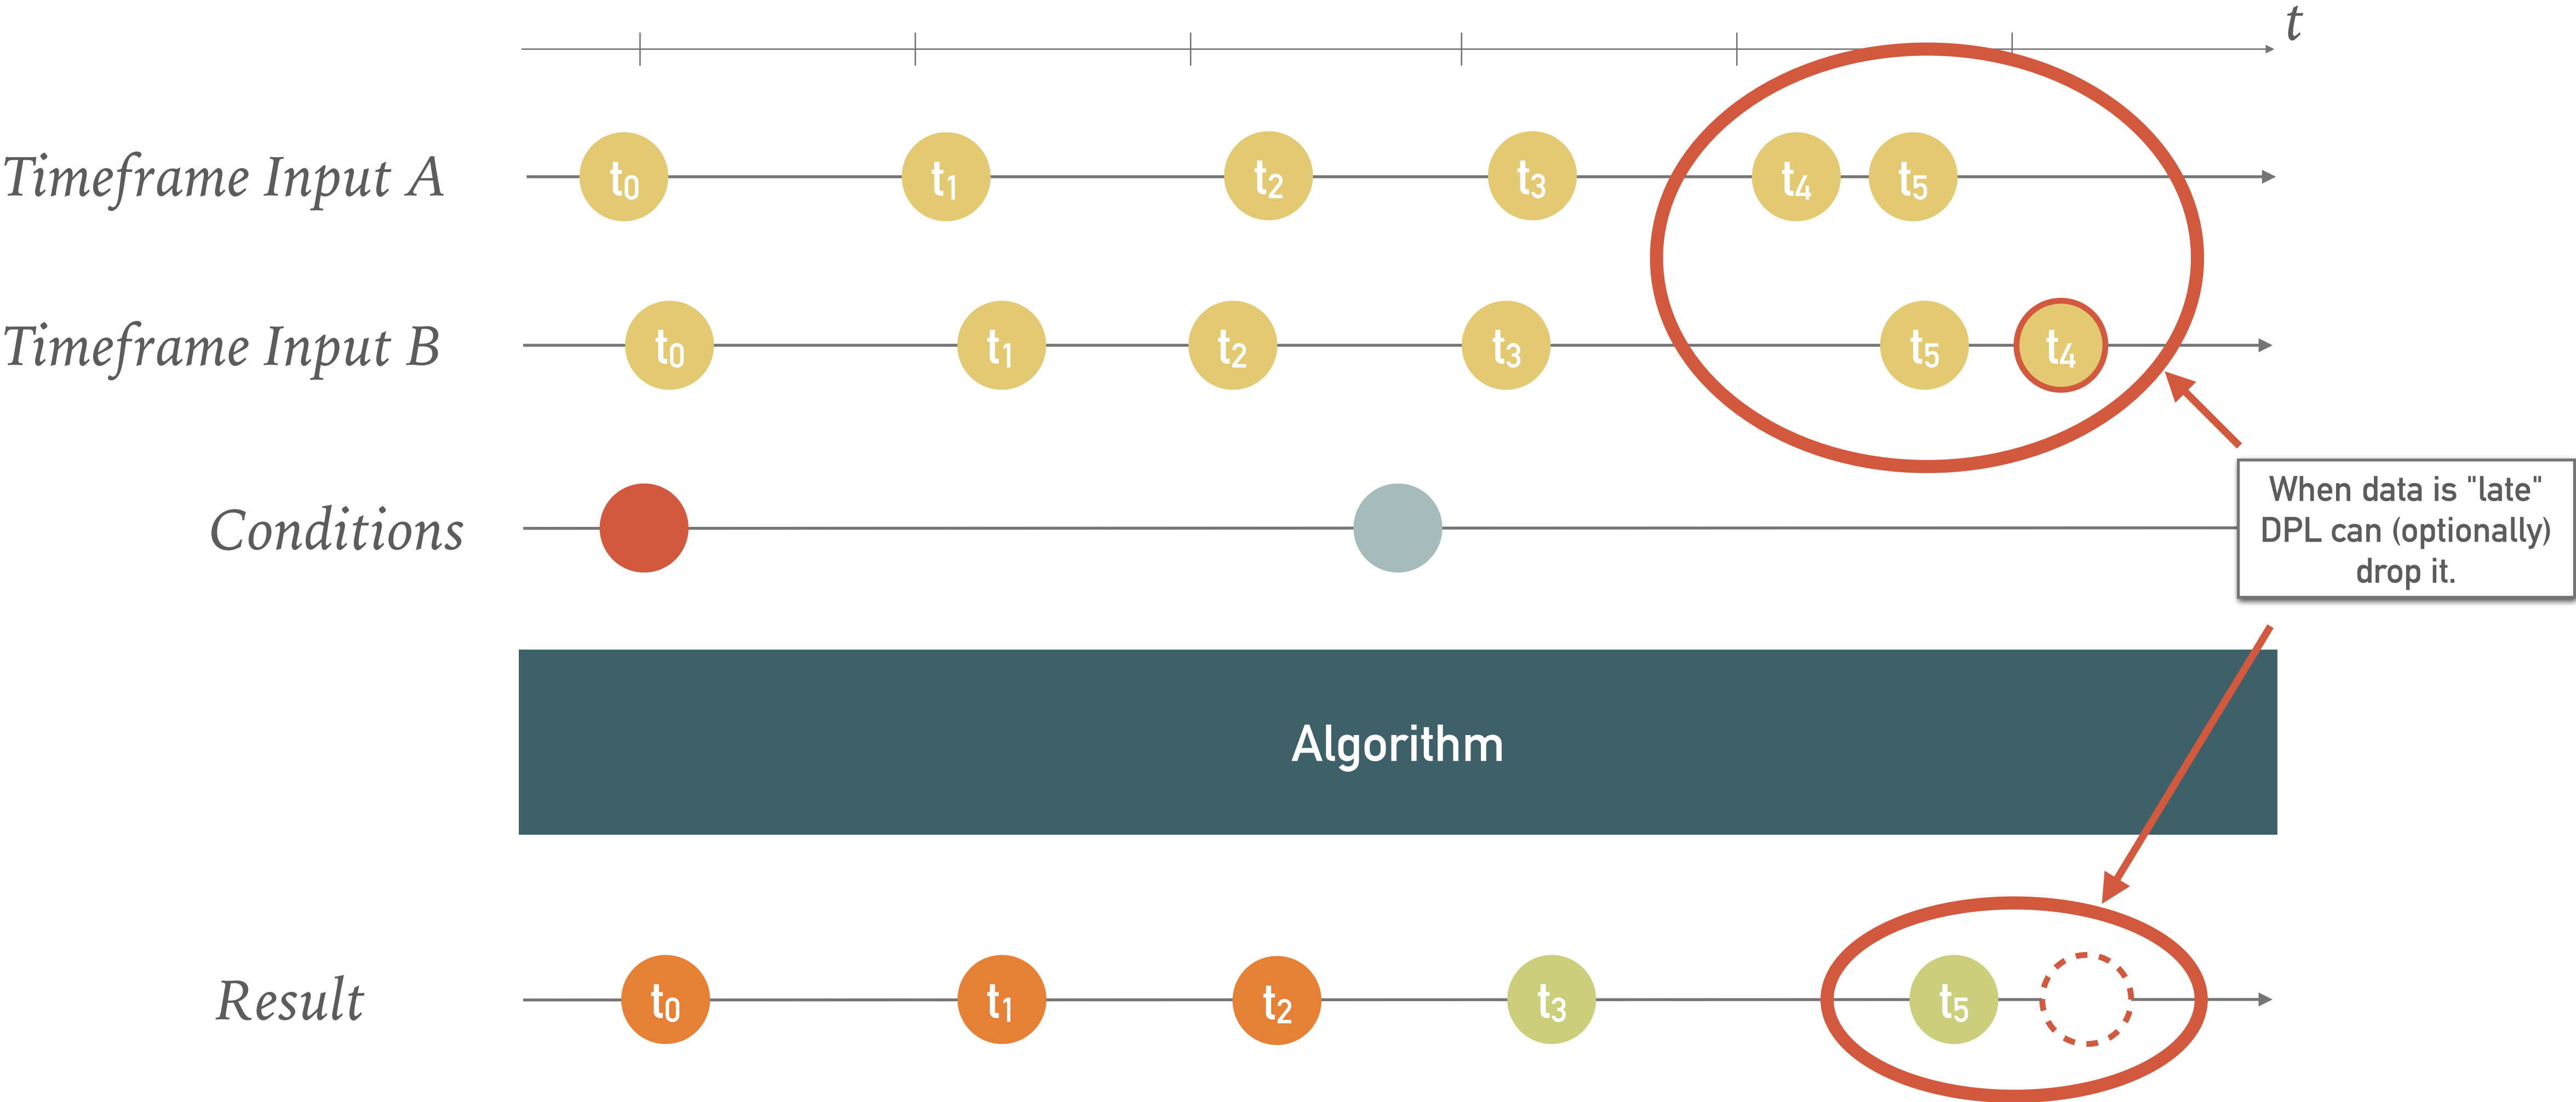


REACTIVE DESIGN




Different inputs can have different lifetimes. E.g. conditions.

REACTIVE DESIGN



DPL as a workflow definition language: support for multiple deployment strategies.

Compiles into a single executable for the laptop user.



Generates DDS configuration for deployment on a analysis facility farm.

```
<?xml version="1.0" encoding="UTF-8" ?>
<topology id="o2-dataflow">
  <decltask id="A">
    <exe reachable="true">../bin/o2DiamondWorkflow --id A ...</exe>
  </decltask>
  <decltask id="B">
    <exe reachable="true">../bin/o2DiamondWorkflow --id B ...</exe>
  </decltask>
  <decltask id="C">
    <exe reachable="true">../bin/o2DiamondWorkflow --id C ...</exe>
  </decltask>
  <decltask id="D">
    <exe reachable="true">../bin/o2DiamondWorkflow --id D ...</exe>
  </decltask>
</topology>
```



..or it generates the configuration to integrate with O2 Control system.

TASK ID (20 TASKS)	CLASS NAME	HOSTNAME	STATUS	STATE
952809ca-e8df-11e8-ace4-a08cfdc880fc	source-1	192.168.05.111	ACTIVE	RUNNING
9527fa39-e8df-11e8-ace4-a08cfdc880fc	step-1	192.168.05.111	ACTIVE	RUNNING
9527ea62-e8df-11e8-ace4-a08cfdc880fc	Dispatcher1	192.168.05.111	ACTIVE	RUNNING
9527d1cb-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask1	192.168.05.111	ACTIVE	RUNNING
9527b950-e8df-11e8-ace4-a08cfdc880fc	source-2	192.168.05.111	ACTIVE	RUNNING
9527ac21-e8df-11e8-ace4-a08cfdc880fc	step-2	192.168.05.111	ACTIVE	RUNNING
95279e10-e8df-11e8-ace4-a08cfdc880fc	sink-2	192.168.05.111	ACTIVE	RUNNING
95278df9-e8df-11e8-ace4-a08cfdc880fc	Dispatcher2	192.168.05.111	ACTIVE	RUNNING
95277e12-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask2	192.168.05.111	ACTIVE	RUNNING
95276cf8-e8df-11e8-ace4-a08cfdc880fc	source-3	192.168.05.111	ACTIVE	RUNNING
95275f23-e8df-11e8-ace4-a08cfdc880fc	step-3	192.168.05.111	ACTIVE	RUNNING
952750a1-e8df-11e8-ace4-a08cfdc880fc	Dispatcher3	192.168.05.111	ACTIVE	RUNNING
952742b4-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask3	192.168.05.111	ACTIVE	RUNNING
95273620-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask-merger	192.168.05.111	ACTIVE	RUNNING
952728b3-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask-checker	192.168.05.111	ACTIVE	RUNNING
95271bc5-e8df-11e8-ace4-a08cfdc880fc	someNumbersTask	192.168.05.111	ACTIVE	RUNNING
95270902-e8df-11e8-ace4-a08cfdc880fc	someNumbersTask-checker	192.168.05.111	ACTIVE	RUNNING
9526f670-e8df-11e8-ace4-a08cfdc880fc	sink-1	192.168.05.111	ACTIVE	RUNNING
9526e183-e8df-11e8-ace4-a08cfdc880fc	sink-3	192.168.05.111	ACTIVE	RUNNING
95267e1e-e8df-11e8-ace4-a08cfdc880fc	dpl-global-binary-file-sink	192.168.05.111	ACTIVE	RUNNING


```
workflow:
[RUNNING] qc-advanced-root
  [RUNNING] source-1
  [RUNNING] step-1
  [RUNNING] Dispatcher1
  [RUNNING] dataSizeTask1
  [RUNNING] source-2
  [RUNNING] step-2
  [RUNNING] sink-2
  [RUNNING] Dispatcher2
  [RUNNING] dataSizeTask2
  [RUNNING] source-3
  [RUNNING] step-3
  [RUNNING] Dispatcher3
  [RUNNING] dataSizeTask3
  [RUNNING] dataSizeTask-merger
  [RUNNING] dataSizeTask-checker
  [RUNNING] someNumbersTask
  [RUNNING] someNumbersTask-checker
  [RUNNING] sink-1
  [RUNNING] sink-3
  [RUNNING] dpl-global-binary-file-sink
```

DPL INTEGRATION: CONTROL

Using DPL will guarantee easy deployment in the synchronous workflow (Teo).

```
teo@pcald15 ~$ coconut env show 9521a76c-e8df-11e8-ace4-a08cfdc880fc -tw
environment id: 9521a76c-e8df-11e8-ace4-a08cfdc880fc
created: 2018-11-15 15:06:01 CET
state: CONFIGURED
```

TASK ID (20 TASKS)	CLASS NAME	HOSTNAME	STATUS	STATE
952809ca-e8df-11e8-ace4-a08cfdc880fc	source-1	192.168.65.111	ACTIVE	CONFIGURED
9527fa39-e8df-11e8-ace4-a08cfdc880fc	step-1	192.168.65.111	ACTIVE	CONFIGURED
9527ea62-e8df-11e8-ace4-a08cfdc880fc	Dispatcher1	192.168.65.111	ACTIVE	CONFIGURED
9527d1cb-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask1	192.168.65.111	ACTIVE	CONFIGURED
9527b950-e8df-11e8-ace4-a08cfdc880fc	source-2	192.168.65.111	ACTIVE	CONFIGURED
9527ac21-e8df-11e8-ace4-a08cfdc880fc	step-2	192.168.65.111	ACTIVE	CONFIGURED
95279e16-e8df-11e8-ace4-a08cfdc880fc	sink-2	192.168.65.111	ACTIVE	CONFIGURED
95278df9-e8df-11e8-ace4-a08cfdc880fc	Dispatcher2	192.168.65.111	ACTIVE	CONFIGURED
95277e12-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask2	192.168.65.111	ACTIVE	CONFIGURED
95276cf8-e8df-11e8-ace4-a08cfdc880fc	source-3	192.168.65.111	ACTIVE	CONFIGURED
95275f23-e8df-11e8-ace4-a08cfdc880fc	step-3	192.168.65.111	ACTIVE	CONFIGURED
952750a1-e8df-11e8-ace4-a08cfdc880fc	Dispatcher3	192.168.65.111	ACTIVE	CONFIGURED
952742b4-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask3	192.168.65.111	ACTIVE	CONFIGURED
95273626-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask-merger	192.168.65.111	ACTIVE	CONFIGURED
952728b3-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask-checker	192.168.65.111	ACTIVE	CONFIGURED
95271bc5-e8df-11e8-ace4-a08cfdc880fc	someNumbersTask	192.168.65.111	ACTIVE	CONFIGURED
95270902-e8df-11e8-ace4-a08cfdc880fc	someNumbersTask-checker	192.168.65.111	ACTIVE	CONFIGURED
9526f670-e8df-11e8-ace4-a08cfdc880fc	sink-1	192.168.65.111	ACTIVE	CONFIGURED
9526e183-e8df-11e8-ace4-a08cfdc880fc	sink-3	192.168.65.111	ACTIVE	CONFIGURED
95267e1e-e8df-11e8-ace4-a08cfdc880fc	dpl-global-binary-file-sink	192.168.65.111	ACTIVE	CONFIGURED

```
workflow:
[CONFIGURED] qc-advanced-root
  [CONFIGURED] source-1 → task 952809ca-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] step-1 → task 9527fa39-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] Dispatcher1 → task 9527ea62-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] dataSizeTask1 → task 9527d1cb-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] source-2 → task 9527b950-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] step-2 → task 9527ac21-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] sink-2 → task 95279e16-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] Dispatcher2 → task 95278df9-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] dataSizeTask2 → task 95277e12-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] source-3 → task 95276cf8-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] step-3 → task 95275f23-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] Dispatcher3 → task 952750a1-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] dataSizeTask3 → task 952742b4-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] dataSizeTask-merger → task 95273626-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] dataSizeTask-checker → task 952728b3-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] someNumbersTask → task 95271bc5-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] someNumbersTask-checker → task 95270902-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] sink-1 → task 9526f670-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] sink-3 → task 9526e183-e8df-11e8-ace4-a08cfdc880fc
  [CONFIGURED] dpl-global-binary-file-sink → task 95267e1e-e8df-11e8-ace4-a08cfdc880fc
```

```
teo@pcald15 ~$ coconut env control 9521a76c-e8df-11e8-ace4-a08cfdc880fc -e START_ACTIVITY
transition complete
environment id: 9521a76c-e8df-11e8-ace4-a08cfdc880fc
state: RUNNING
```

```
teo@pcald15 ~$ coconut env show 9521a76c-e8df-11e8-ace4-a08cfdc880fc -tw
environment id: 9521a76c-e8df-11e8-ace4-a08cfdc880fc
created: 2018-11-15 15:06:01 CET
state: RUNNING
```

TASK ID (20 TASKS)	CLASS NAME	HOSTNAME	STATUS	STATE
952809ca-e8df-11e8-ace4-a08cfdc880fc	source-1	192.168.65.111	ACTIVE	RUNNING
9527fa39-e8df-11e8-ace4-a08cfdc880fc	step-1	192.168.65.111	ACTIVE	RUNNING
9527ea62-e8df-11e8-ace4-a08cfdc880fc	Dispatcher1	192.168.65.111	ACTIVE	RUNNING
9527d1cb-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask1	192.168.65.111	ACTIVE	RUNNING
9527b950-e8df-11e8-ace4-a08cfdc880fc	source-2	192.168.65.111	ACTIVE	RUNNING
9527ac21-e8df-11e8-ace4-a08cfdc880fc	step-2	192.168.65.111	ACTIVE	RUNNING
95279e16-e8df-11e8-ace4-a08cfdc880fc	sink-2	192.168.65.111	ACTIVE	RUNNING
95278df9-e8df-11e8-ace4-a08cfdc880fc	Dispatcher2	192.168.65.111	ACTIVE	RUNNING
95277e12-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask2	192.168.65.111	ACTIVE	RUNNING
95276cf8-e8df-11e8-ace4-a08cfdc880fc	source-3	192.168.65.111	ACTIVE	RUNNING
95275f23-e8df-11e8-ace4-a08cfdc880fc	step-3	192.168.65.111	ACTIVE	RUNNING
952750a1-e8df-11e8-ace4-a08cfdc880fc	Dispatcher3	192.168.65.111	ACTIVE	RUNNING
952742b4-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask3	192.168.65.111	ACTIVE	RUNNING
95273626-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask-merger	192.168.65.111	ACTIVE	RUNNING
952728b3-e8df-11e8-ace4-a08cfdc880fc	dataSizeTask-checker	192.168.65.111	ACTIVE	RUNNING
95271bc5-e8df-11e8-ace4-a08cfdc880fc	someNumbersTask	192.168.65.111	ACTIVE	RUNNING
95270902-e8df-11e8-ace4-a08cfdc880fc	someNumbersTask-checker	192.168.65.111	ACTIVE	RUNNING
9526f670-e8df-11e8-ace4-a08cfdc880fc	sink-1	192.168.65.111	ACTIVE	RUNNING
9526e183-e8df-11e8-ace4-a08cfdc880fc	sink-3	192.168.65.111	ACTIVE	RUNNING
95267e1e-e8df-11e8-ace4-a08cfdc880fc	dpl-global-binary-file-sink	192.168.65.111	ACTIVE	RUNNING

```
workflow:
[RUNNING] qc-advanced-root
  [RUNNING] source-1 → task 952809ca-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] step-1 → task 9527fa39-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] Dispatcher1 → task 9527ea62-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] dataSizeTask1 → task 9527d1cb-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] source-2 → task 9527b950-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] step-2 → task 9527ac21-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] sink-2 → task 95279e16-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] Dispatcher2 → task 95278df9-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] dataSizeTask2 → task 95277e12-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] source-3 → task 95276cf8-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] step-3 → task 95275f23-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] Dispatcher3 → task 952750a1-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] dataSizeTask3 → task 952742b4-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] dataSizeTask-merger → task 95273626-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] dataSizeTask-checker → task 952728b3-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] someNumbersTask → task 95271bc5-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] someNumbersTask-checker → task 95270902-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] sink-1 → task 9526f670-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] sink-3 → task 9526e183-e8df-11e8-ace4-a08cfdc880fc
  [RUNNING] dpl-global-binary-file-sink → task 95267e1e-e8df-11e8-ace4-a08cfdc880fc
```

INTEGRATION WITH DIFFERENT SUBSYSTEMS

Monitoring infrastructure

Sending metrics to the Debug GUI and to the Monitoring infrastructure is now done with the same code:

```
ctx.services().get<Monitoring>().send({ 1, "metric-name" });
```

Deployment target will pick up the correct backend.

InfoLogger

Log messages (e.g. those created via FairLogger LOG() macro) will be redirected to the specified InfoLogger back-end (or the Debug GUI).

O2 Control & Grid (preliminary)

DPL generates configuration templates which are suitable to be imported in the O2 Control system or in DDS (e.g. for Grid deployments).

Configuration Database (coming soon)

DPL already provides a way to specify configurable options. At the moment these are only mapped to command-line arguments but it's in the plan to fully expose them via FairMQProgOptions and integrate DPL with the Configuration Database.

MISC RECENT DEVELOPMENTS

Special Input / Output types

Not all data is part of the main data-flow (e.g. the timeframes) but we will have timers (e.g. to trigger histograms snapshots in QA), condition objects (which will come from / be stored in CCDB via a REST API call). Out-of-band input / output types are now supported and effort is ongoing to integrate them with different backends.

Support for Apache Arrow & RDataFrame

We now have initial support for Apache Arrow backed messages, which paves the way to integration with external tools like Pandas and TensorFlow, or even ROOT, via RDataFrame.

COMING UP NEXT

PR #1483 is paving the way for a major refactoring which will provide:

- More compact syntax to specify inputs

From:

```
{ InputSpec{ "clusters", "TPC", "CLUSTERS", 0 },  
  InputSpec{ "tracks", "TPC", "TRACKS", 0 }, ... }
```

to:

```
select ("clusters:TPC/CLUSTERS/0;tracks:TPC/TRACKS/0")
```

- Wildcards when selecting input data:

```
select ("clusters:$1/CLUSTERS/0;tracks:$1/TRACKS/0")
```

- Different time granularities

```
select ("test:A/D1/0%Timeframe;test:A/D2/0%Run")
```

COMING UP AFTER

Shared Services

DataProcessorSpec API will be extended to request shared services (e.g. Geometry, Magnetic Field) which will be shared between all instances running on the same NUMA-domain.

Workflow merging

Ability to merge separate workflows into a larger ones and run it / generate the appropriate deployment configuration. E.g. digitization + reconstruction.

Resource limiting

Resources like CPUs and memory will be modelled as tokens exchanged by various components so that only a limited number of them can be active at the same time. Useful in particular for Grid workflows where we might have a very small number of cores per job.

MORE INFO

Design document and cookbook

Available from our doxygen web pages <https://aliceo2group.github.io/AliceO2>

Examples

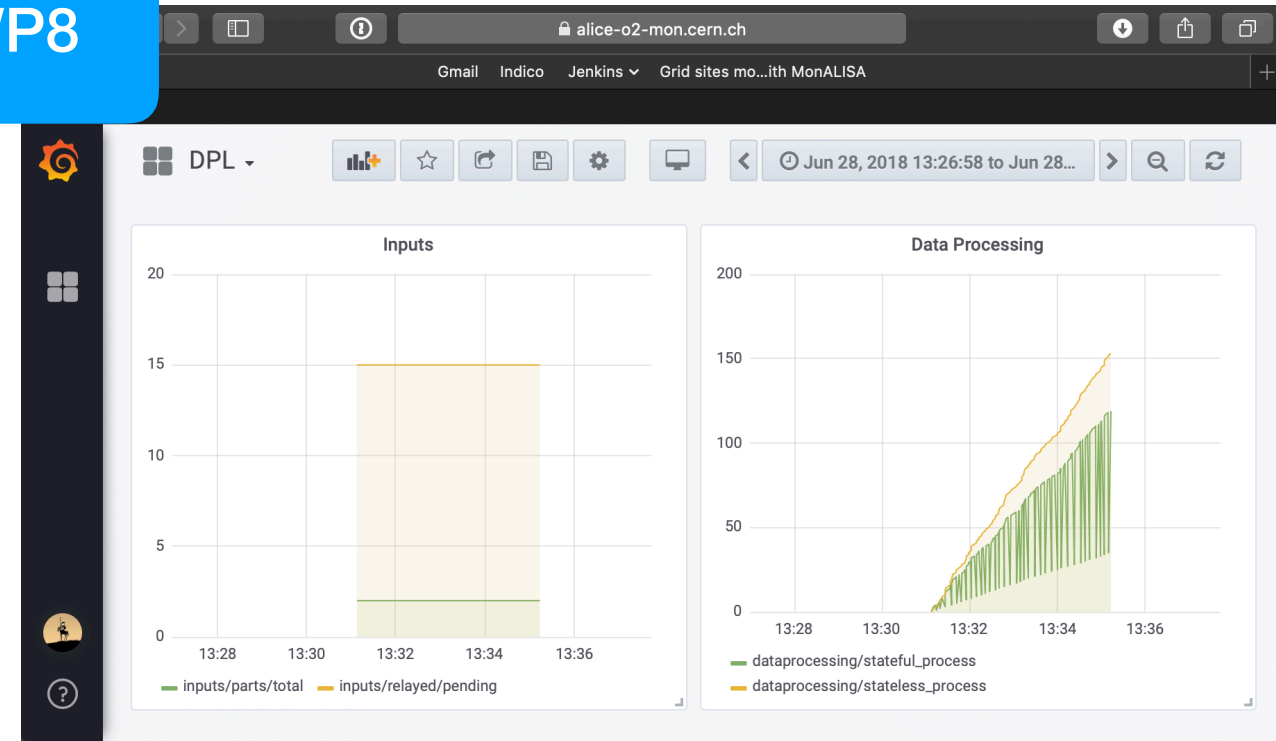
AliceO2/Framework/Core/test and *AliceO2/Framework/TestWorkflows*

WP4 mailing egroup

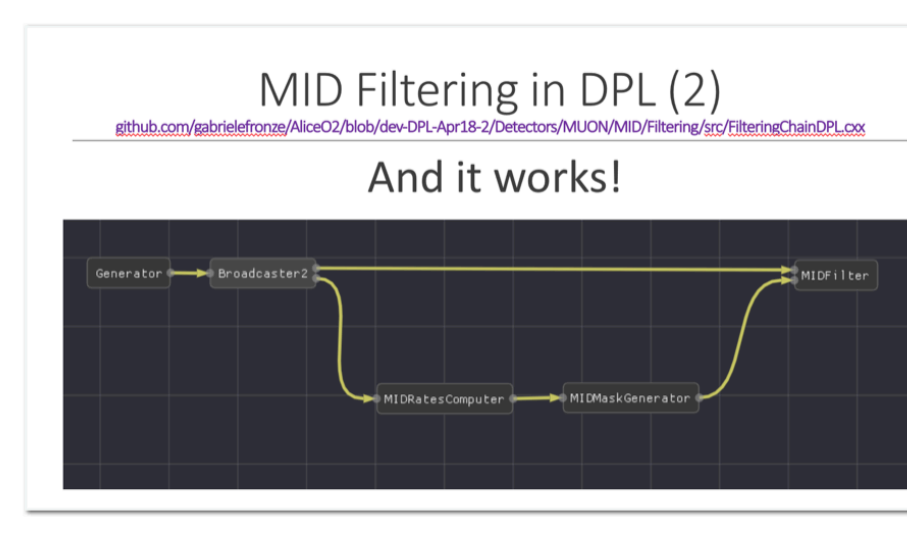
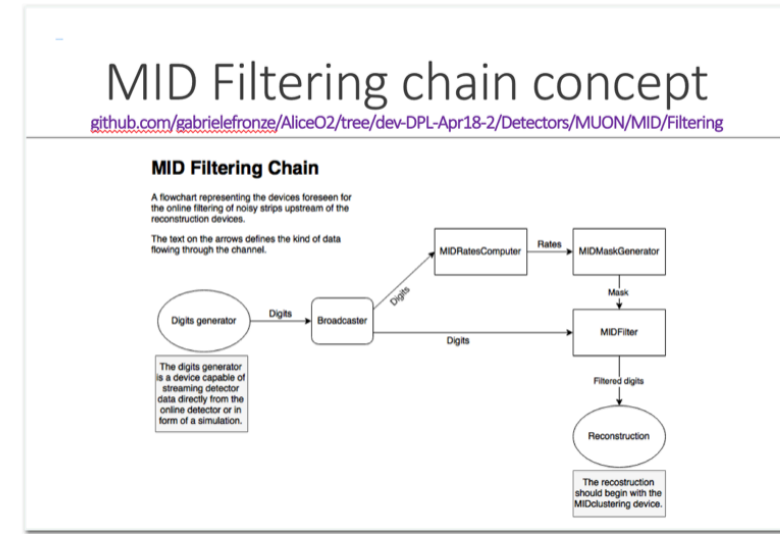
alice-o2-wp4@cern.ch

DPL AS AN INTEGRATION PLATFORM FOR O2

WP8



O2 Monitoring and InfoLogger integration



MID Filtering Chain (kudos Gabriele Fronzè)

```

$> tcoincalis coconut env show 9521a76c-e8df-11e8-ace4-a08cfdc880fc -tw
environment id: 9521a76c-e8df-11e8-ace4-a08cfdc880fc
created: 2018-11-15 15:06:01 CET
state: RUNNING

TASK ID (20 TASKS) | CLASS NAME | HOSTNAME | STATUS | STATE
-----|-----|-----|-----|-----
952809ca-e8df-11e8-ace4-a08cfdc880fc | source-1 | 192.168.65.111 | ACTIVE | RUNNING
9527fa39-e8df-11e8-ace4-a08cfdc880fc | step-1 | 192.168.65.111 | ACTIVE | RUNNING
9527a652-e8df-11e8-ace4-a08cfdc880fc | Dispatcher1 | 192.168.65.111 | ACTIVE | RUNNING
9527d1cb-e8df-11e8-ace4-a08cfdc880fc | dataSizeTask1 | 192.168.65.111 | ACTIVE | RUNNING
9527b950-e8df-11e8-ace4-a08cfdc880fc | source-2 | 192.168.65.111 | ACTIVE | RUNNING
9527ac21-e8df-11e8-ace4-a08cfdc880fc | step-2 | 192.168.65.111 | ACTIVE | RUNNING
95279e16-e8df-11e8-ace4-a08cfdc880fc | sink-2 | 192.168.65.111 | ACTIVE | RUNNING
95278df9-e8df-11e8-ace4-a08cfdc880fc | Dispatcher2 | 192.168.65.111 | ACTIVE | RUNNING
95277e12-e8df-11e8-ace4-a08cfdc880fc | dataSizeTask2 | 192.168.65.111 | ACTIVE | RUNNING
9527cf8-e8df-11e8-ace4-a08cfdc880fc | source-3 | 192.168.65.111 | ACTIVE | RUNNING
95275f23-e8df-11e8-ace4-a08cfdc880fc | step-3 | 192.168.65.111 | ACTIVE | RUNNING
952750a1-e8df-11e8-ace4-a08cfdc880fc | Dispatcher3 | 192.168.65.111 | ACTIVE | RUNNING
952742b4-e8df-11e8-ace4-a08cfdc880fc | dataSizeTask3 | 192.168.65.111 | ACTIVE | RUNNING
95273626-e8df-11e8-ace4-a08cfdc880fc | dataSizeTask-merger | 192.168.65.111 | ACTIVE | RUNNING
952728b3-e8df-11e8-ace4-a08cfdc880fc | dataSizeTask-checker | 192.168.65.111 | ACTIVE | RUNNING
95271bc5-e8df-11e8-ace4-a08cfdc880fc | someNumbersTask | 192.168.65.111 | ACTIVE | RUNNING
95270902-e8df-11e8-ace4-a08cfdc880fc | someNumbersTask-checker | 192.168.65.111 | ACTIVE | RUNNING
9526f670-e8df-11e8-ace4-a08cfdc880fc | sink-1 | 192.168.65.111 | ACTIVE | RUNNING
9526e183-e8df-11e8-ace4-a08cfdc880fc | sink-3 | 192.168.65.111 | ACTIVE | RUNNING
95267e1e-e8df-11e8-ace4-a08cfdc880fc | dpl-global-binary-file-sink | 192.168.65.111 | ACTIVE | RUNNING

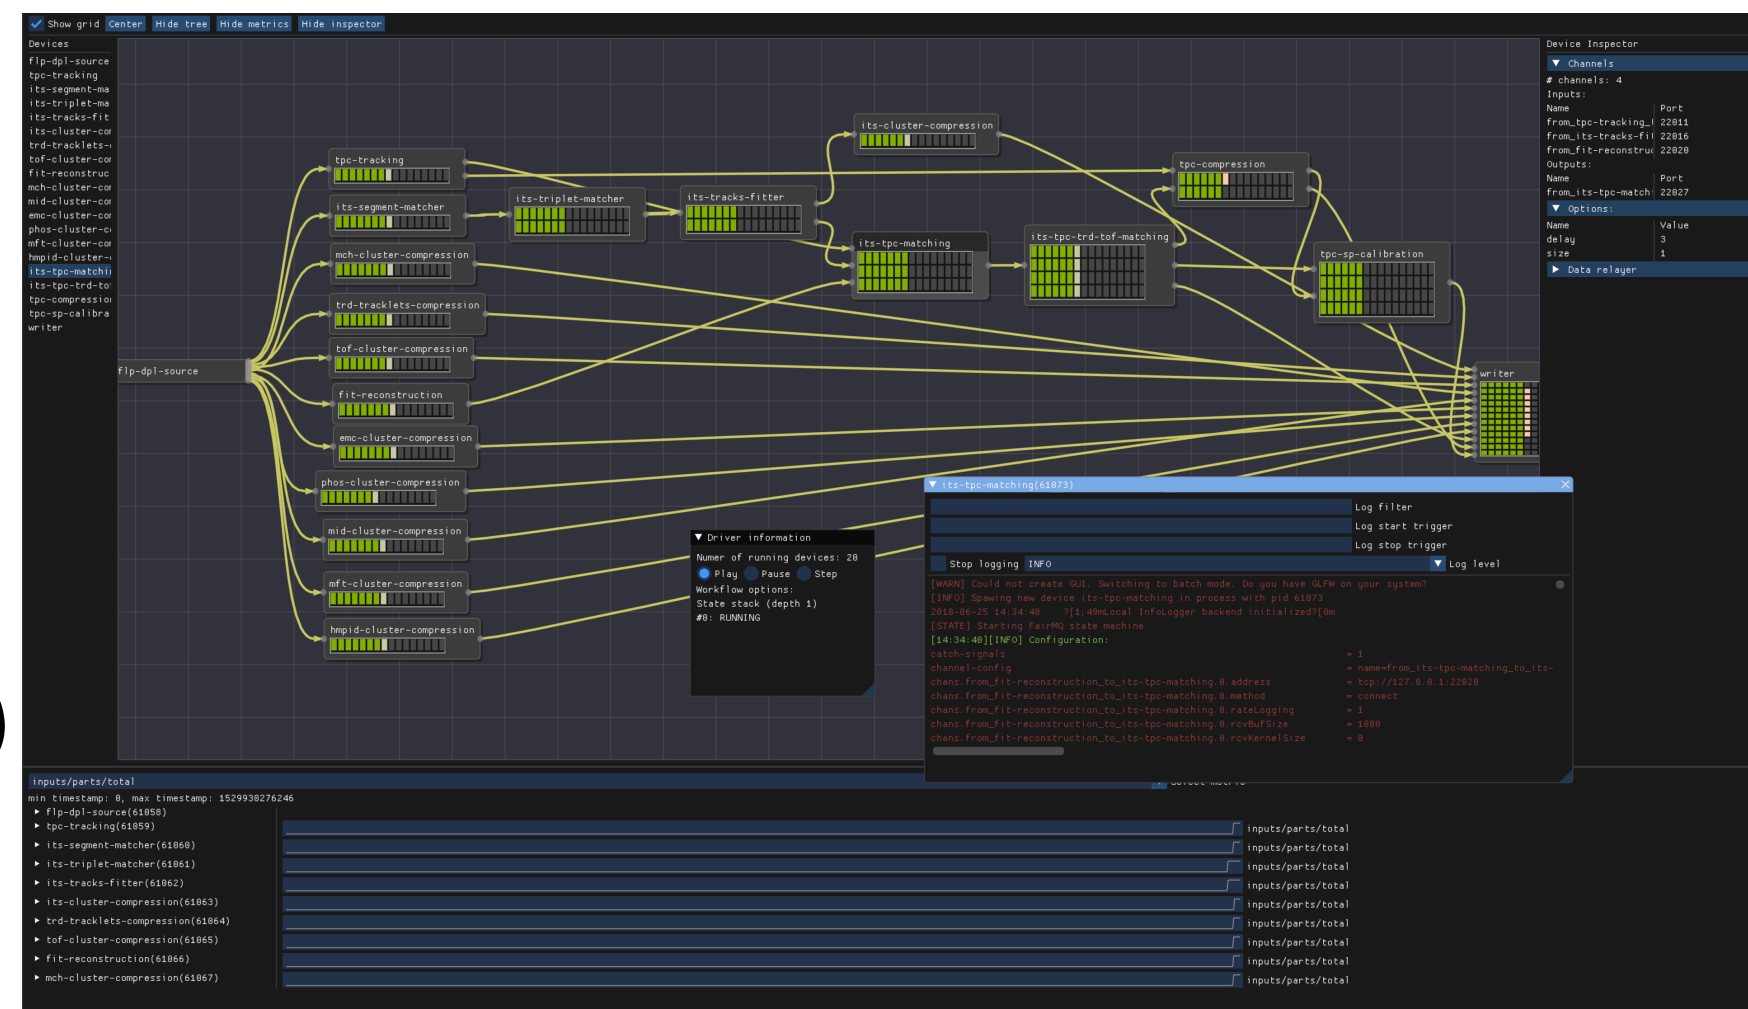
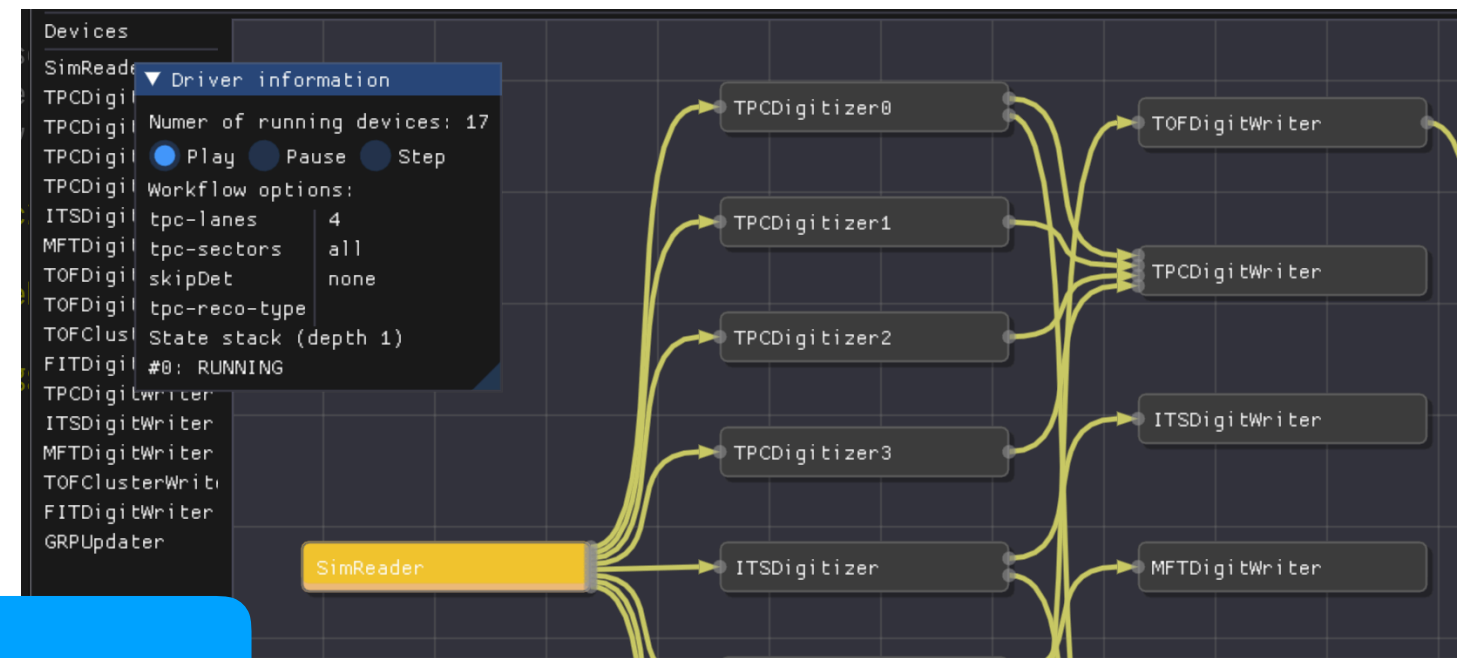
workflow:
[RUNNING] qc-advanced-root
├── [RUNNING] source-1
├── [RUNNING] step-1
├── [RUNNING] Dispatcher1
├── [RUNNING] dataSizeTask1
├── [RUNNING] source-2
├── [RUNNING] step-2
├── [RUNNING] sink-2
├── [RUNNING] Dispatcher2
├── [RUNNING] dataSizeTask2
├── [RUNNING] source-3
├── [RUNNING] step-3
├── [RUNNING] Dispatcher3
├── [RUNNING] dataSizeTask3
├── [RUNNING] dataSizeTask-merger
├── [RUNNING] dataSizeTask-checker
├── [RUNNING] someNumbersTask
├── [RUNNING] someNumbersTask-checker
├── [RUNNING] sink-1
├── [RUNNING] sink-3
└── [RUNNING] dpl-global-binary-file-sink
    
```

O2 Control Integration

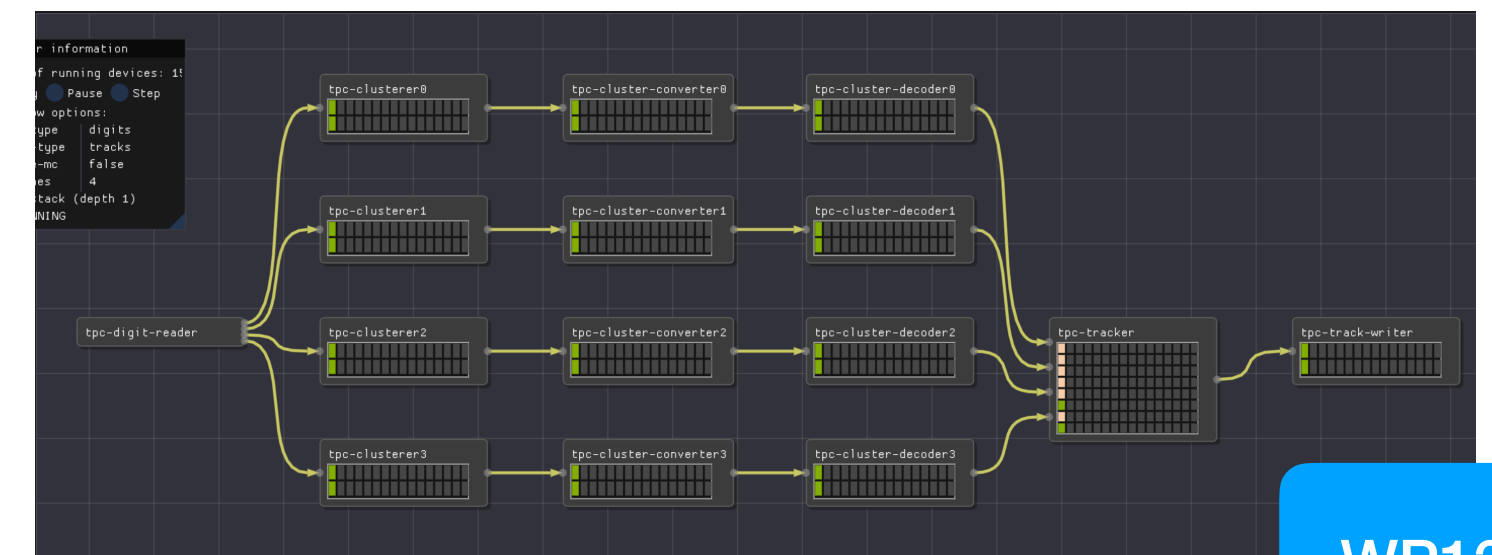
WP8

WP12

Digitization in DPL (see Sandro's talk)



ALICE Framework showcased at CHEP2018 Sofia. Establishing the Data Processing Layer (DPL) as integration framework for ALICE data-processing needs.



TPC reconstruction prototype (see Jens's talk)

WP13

WP7

DataSampling using DPL (kudos Piotr)

